

Uncertainty Quantified Computational Analysis of the Energetics of Virus Capsid Assembly

Nathan L Clement,^{†,‡} Muhibur Rasheed,^{†,‡} and Chandrajit L Bajaj^{*,†}

Department of Computer Science, The University of Texas at Austin, Austin, TX 78712

E-mail: bajaj@cs.utexas.edu

*To whom correspondence should be addressed

[†]University of Texas

[‡]Both authors contributed equally to this work

Abstract

Most of the existing research in assembly pathway prediction/analysis of virus capsids makes the simplifying assumption that the configuration of the intermediate states can be extracted directly from the final configuration of the entire capsid. This assumption does not take into account the conformational changes of the constituent proteins as well as minor changes to the binding interfaces that continues throughout the assembly process until stabilization. This paper presents a statistical-ensemble based approach which provides sufficient samples of the configurational space for each monomer and the relative local orientation between monomers, to capture the uncertainties in their binding and conformations. Furthermore, instead of using larger capsomers (trimers, pentamers) as building blocks, we allow all possible sub-assemblies to bind in all possible combinations. We represent this assembly graph in two different ways. First, we use the Wilcoxon signed rank measure to compare the distributions of binding free energy computed on the sampled conformations to predict likely pathways. Second, we represent chemical equilibrium aspects of the transitions as a Bayesian Factor graph where both associations and dissociations are modeled based on concentrations and the binding free energies. Results from both of these experiments showed significant departure from those one would obtain if only the static configurations of the proteins were considered. Hence, we establish the importance of an uncertainty-aware protocol for pathway analysis, and provide a statistical framework as an important first step towards assembly pathway prediction with high statistical confidence.

Introduction

Viruses are the smallest living organisms on earth, possessing only a minimal genome which can transcribe only a few proteins. Even with such limited resources, viruses exhibit a remarkable ability to not only survive, but parasitically multiply with great efficiency. One phase of their chemical proliferation that is relatively unexplained, is the spontaneous assembly/disassembly of hundreds of (capsid) proteins that come together to form the capsid (shell) that encases the viral nucleic acid genetic material. Researchers continue to analyze this remarkable process from different perspectives, and also aim to use the insights discovered in designing nano-scale cages and shells for drug delivery.¹⁻⁴ In this article, we present a statistical methodology to analyze the viral assembly process from a free energy perspective. We consider in particular the case of the Nudaurelia Capensis virus (PDBID:1OHF⁵). While others have taken a similar perspective before (e.g.⁶⁻¹⁰), we are the first to consider positional and conformational uncertainties of the protein structure and their propagated influence on the configurational energetics and binding affinity calculations. This methodology then allows us to infer energetically favorable viral capsomer configurations together with improved statistical confidence.

Research at understanding the assembled arrangement of capsid proteins in a viral capsid builds upon the work of Caspar and Klug (C-K).¹¹ C-K characterize the symmetric organization of proteins in a spherical viral capsid, building upon the mathematical foundations of spherical tilings given by Goldberg.¹² C-K show that the combinatorial arrangement of the capsid proteins can be characterized using simple triangular tiles that cover an icosahedron. The basic concept is to unfold an icosahedron, which has 20 regular triangular faces (and 12 vertices of 5-fold rotational symmetry), onto a regular hexagonal grid (geometric dual to a triangular grid). By controlling the scale of the icosahedral triangles with respect to the dual-grid-triangles, each icosahedral triangle gets covered by T small *grid-triangles*, such that $T = h^2 + k^2 + hk$ where h and k are integers. C-K suggested that each such grid-triangle consists of 3 proteins in a cyclic configuration (forming a trimer). Additionally, there are exactly 12 locations where 5 such small triangles meet at a vertex (vertices of the icosahedron), and $10(T - 1)$ locations where 6 such triangles meet. If we only

consider the viral protein (of the trimer) which is closest to these 5-fold and 6-fold locations, then they can be thought to produce pentamers (pentons) and hexamers (hexons). Essentially, the entire capsid can be thought of as $20T$ trimers; or as 12 pentamers and $10(T - 1)$ hexamers. Finally, their idea, called “quasi-equivalence,” states that all the viral protein chains that form a capsid are identical and have (quasi-) equivalent interfaces—essentially, all proteins are involved in the same number of interactions at similar binding sites.

Recent published work by Janner¹³ and Keef and Twarock¹⁴ have shown that other types of aperiodic arrangements involving pentamers or other types of subassemblies are also possible. Other work, for example by Pawley,¹⁵ shows that several other symmetry classes also permit decomposition into symmetric subassemblies, and Rasheed *et al.*¹⁶ proved necessary and sufficient conditions for such subassemblies to be possible. Brooks *et al.*, in a series of papers, have characterized the geometric conditions for symmetric capsids, provided methods to measure how much a specific capsid conforms to the concept of quasi-equivalence,^{17–19} and hence how amenable it is to coarse-grained dynamics analysis as described below. Furthermore, they present a simple classification which characterizes variations of hexamers within a capsid.²⁰

Many of the researchers working on predicting, analyzing, and/or simulating capsid assembly have taken either a set of trimers, or a set of pentamers+hexamers, as the building blocks of assembly. For instance, Rapaport *et al.*¹⁰ performed a coarse dynamics simulation where trimers were used as building blocks. The shape of the trimers were modeled using a collection of large balls, and the interactions were modeled by proximity of some small balls placed strategically along the binding region. While this is a very simplistic model, it was successful in showing that such shape and binding site conditions are sufficient to drive self-assembly. In their more recent work (e.g.^{21,22}), the model was updated to include more complex energetics, and single proteins (monomers) were used as building blocks, instead of trimers.

Hagan *et al.*^{6,23} applied Brownian dynamics simulation with a simplified force field. They modeled each capsomer (which can also be a monomer) using a single bead model. Based on prior knowledge about the arrangement of such beads on the capsid, they parametrized each bead

based on the angles between each pair of their neighbors, and designed a binding affinity function which allowed binding at specific orientations. The objective of the study was to gain insight from exploring the energy landscape of the assembly, and to identify kinetic traps, analyze the rate of assembly etc. This concept is similar to the ‘local-rules’ introduced by Berger *et al.*,^{24,25} which has been adopted by other groups,^{26,27} for kinetics and dynamics analysis of capsids. A discussion contrasting the block-like beads used by Rapaport *et al.* with shape-driven assembly, and the ones used by Hagan *et al.* with neighborhood-driven assembly can be found in.⁷ Bona and Sitharam also considered a bead-like model;²⁸ however, they expressed the interaction between beads using geometric stability conditions and predicted likelihood of binding based on the simplicity of solving the resultant geometric constraints system.²⁹

Unlike the dynamics-based analysis techniques described above, Zlotnick applied statistical thermodynamics law of mass action to relate the concentrations of the constituents and the product of a binding with the binding free energy.⁸ Using pentameric building blocks, he enumerated all unique compositions of one or more pentamers (each arranged exactly as it would be if the entire capsid was formed). This technique, and several following publications,^{30,31} revealed various aspects of assembly for different viruses, including rates of assembly, effect of nucleation, detection of possible kinetic traps etc. It also provided a simple tool to predict the effect of changing environment parameters and or presence of other molecules, which can be applied to measure yields under different conditions, designing conditions amenable to specific assemblies, etc.^{9,32}

This paper presents an approach to score and rank conformational ensembles of capsid protein capsomers and capsid subassemblies based on a new configurational sampling and energy analysis approach. The sampled configurations of capsid subassemblies represents the various potential intermediate states of a fully assembled viral capsid. In other words, we recognize that the tertiary structure (fold) of individual subunits as well as binding contacts between subunits may evolve over the span of the entire assembly process, and moreover, may exist in slightly different configurations for the same subassemblies. The presence of such uncertainties imply that any binding free energy computed solely based on the structure and interfaces that exist in the final matured state

of the capsid is not always accurate. Similar uncertainty quantification and uncertainty propagation methods have recently been used for single molecule models^{33,34} but not for combinatorial arrangements of viral capsid proteins in various capsomeric states.

In our approach, given prior knowledge (in the form of statistical distributions) of the nature of uncertainty, we can provide additional theoretical upper bounds on the distributional moments^{35–37} for different properties of viral capsomers, and other quantities of interest or QOI (e.g. the binding free energy). See for e.g.³⁴ for such Azuma-Hoeffding bounds applied to molecular modeling with atomistic positional uncertainty captured by B-factors. Additionally, if the space of configurations is sampled such that low-discrepancy (and also low dispersion) is achieved, then a probability distribution of the QOI can be approximated with bounded error.^{38,39} Such estimation of binding free energies, i.e. as distributions instead of single values, makes it possible to sample energy landscapes through various configurational ensembles, and analyze binding pathways in an efficient and robust way.³⁴ We apply our efficient low-discrepancy product space sampling technique reported in⁴⁰ to generate such low-discrepancy sampled ensembles of viral capsomers. The configuration space for any subassembly of the capsid is a product space of the backbone torsion angles between relatively rigid domains, as well as 3D affine transformations between each pair of neighbors.

We also use the Wilcoxon sign rank test⁴¹ to compare the computed sampled energy distribution of various capsomers. Using this we rank all possible transitions from any given subassembly, as a step towards stochastically predicting and analyzing stable assembly pathways, and with quantified uncertainty. Finally, we use the distributions of free energy calculations to provide an analysis of the relative concentrations of intermediate subassemblies as the capsid is being formed.

Materials and Methods

One of the major goals of this work is to develop a method for viral self-assembly pathway analysis with statistical guarantees. We consider assembly from an equilibrium perspective, where,

given prior knowledge of the final assembled structure, we can uniquely determine the possible subassemblies of different sizes and the possible ways they can be associated/disassociated. This assumption, that the binding free energy of the association governs the success and yield of the reaction, is similar to the work of Zlotnick *et al.*,^{8,9,30–32} but applies a more robust estimation of the binding energy under an uncertainty quantification framework. Additionally, we consider all possible assembly pathways starting from monomers, instead of assuming that trimers, pentamers etc. are the basic building blocks.

The overall methodology for this research is as follows. First, we identify all unique interfaces and unique subassemblies of specific sizes (where “size” is defined as the number of constituent monomers) of a given virus. Second, we sample the space of configurations for each of these subassemblies with restricted range of motion to generate an ensemble of structures in an attempt to capture the uncertainty (flexibility, random perturbations etc.) of the structure of the subassembly. Then, we compute the free energy of each sample of each subassembly to generate a distribution of the energy. Finally, we use these distributions of energies (instead of the traditional single value) to compare the stabilities of the subassemblies, compute distributions of binding free energies between subassemblies, predict likely transition pathways from one subassembly to another, etc. In the following subsections, we discuss each of these in detail.

Unique Subassemblies and Transitions

Analysis of self-assembly focusing on only a predetermined set of pathways, fails to take into account subassemblies caught into energy traps, which nonetheless, may be part of a pathway which is globally favorable to the capsid as a whole.^{6,23} For this reason, we have sought to implement an exhaustive approach. We consider all possible unique subassemblies and all possible ways they can come together.

To begin, we consider each chain to be unique. Even though the chains have the same primary structure, in most cases they exhibit minor differences in their tertiary configuration and hence it is preferable to consider each of them as unique, especially when computing binding free energies.

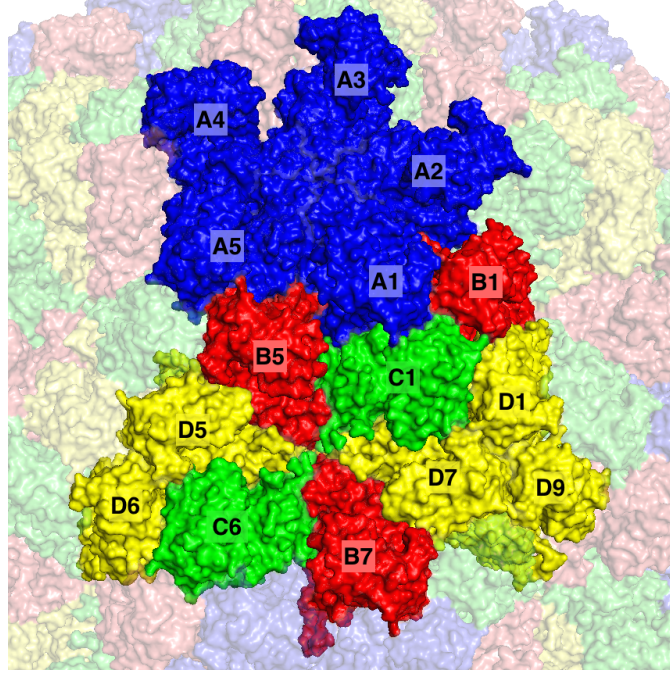


Figure 1: A portion of the capsid of Nudaurelia Capensis virus (PDBID:1OHF) capsid. Labels on the capsid shows individual monomeric capsid proteins of different types A/B/C/D and at different locations and forming different local subassemblies. For instance the capsomer A1-A2-A3-A4-A5 form a pentameric configuration and contains four 5-fold interfaces of the same type. Similarly A1-B1-C1 and D1-D7-D9 are two trimers, but involve slightly different interfaces.

For subassemblies involving two or more monomers, we consider them to be equivalent if and only if all three following conditions (evaluated in this order) are met: 1) they have the same number of monomers of each type, 2) they have the same number of symmetric interfaces of each type, and 3) when the entire subassemblies are aligned, the RMSD is less than a threshold.

For example, the Nudaurelia Capensis virus (PDBID:1OHF) has 240 proteins on its capsid, and has four unique monomers: A, B, C, and D. It has several unique symmetric interfaces, each appearing multiple times on the capsid. In Figure 1, we show a portion of capsid where each of these unique interface types are present at least once. For example, 5-fold between A1-A2, A2-A3, etc.; 6-fold between C1-B5, B5-D5, C1-D7, etc; 3-fold between A1-B1, A1-C1, B1-C1, D1-D7 etc.; and 2-fold between C1-D1, A1-B5 etc. According to our criterion, A1-B1 and A5-B5 are equivalent to each other but not equivalent to A1-B5 (violates criterion 2). A4-A5-A1-B5 and A5-A1-A2-B1 are equivalent, but A1-A2-A3-B1 are not (violates criterion 3).

We select a set of subassemblies such that no member of the set is equivalent to any other member. We ended up with 985 subassemblies involving up to 6 monomers. This set includes some of the more distinct capsomers of this capsid: the trimers (A1-B1-C1) and (D1-D7-D9), the pentamer (A1-A2-A3-A4-A5), and the hexamer (B5-C1-D7-B7-C6-D5). For this work, the number of subunits in any given subassembly was limited to 6.

We consider a transition from subassembly **P** to subassembly **Q** feasible if it is possible to add one monomer to **P** to make it equivalent to **Q**. For example, for the case shown in Figure 1, (A1-B1-C1) is reachable from (A1-B1), (B1-C1), and (A1-C1).

Low-Discrepancy Sampling of Subassemblies

As mentioned before, instead of considering a static model for a subassembly, we are interested in modeling its structure as a distribution of possible structures which have minor differences, but represent the same state. One way to think of this is to consider an energy well that contains the specific subassembly and many others which are just slightly different—in such case, one should not focus on only one of them to characterize the well, but should consider the entire distribution. In this regard, we are interested in both small changes inside subunit conformation (the natural shift in structure of the protein backbone) and slight perturbations of the interface. Now we describe a parameterization of these spaces.

Configuration space of a subassembly

Internal (flexible) DOFs While in principle backbone torsional angles are all relevant when considering internal flexibility of a protein, for the sake of computational tractability (especially for the multitude of subassemblies we have), we applied a coarse-grained approach based on domain decomposition. Essentially, we limited the sampling space to flexible backbone torsion angles between relatively rigid subdomains. To determine the set of flexible backbone torsion angles, we used HingeProt⁴² to identify *hinge* residues for each subunit, P_X , designating the corresponding ϕ and ψ internal torsion angles of each residue as flexible (i.e. if there were r hinge residues,

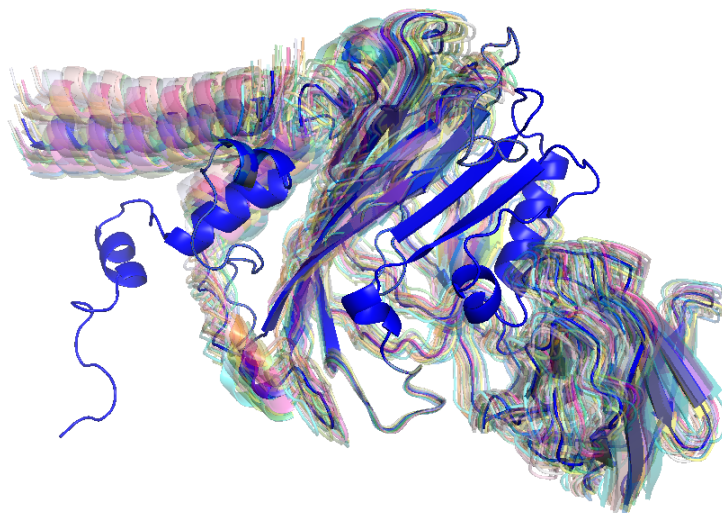


Figure 2: A few sample configurations for monomer chain A, generated by sampling backbone torsion angles, shown using transparent rendering. The original configuration is rendered opaque.

there were a total of $2r$ rotatable bonds). This results in a configurational space equivalent to \mathbb{R}^{2r} . Figure 2 shows an example of sampling such a space.

External DOFs We parametrize the space of local affine perturbations of each pairwise interface between every pair of monomers in a subassembly using 6-DOFs, defined by three Euler angle twists and three translational shifts. Hence, for a subassembly with t pairwise interfaces (edges in a graph like the one in Figure 3), the space is equivalent to \mathbb{R}^{6t}

Sampling

Recall that we want to estimate the distribution of free energy over configurations in a local neighborhood of a given subassembly. Computing such distributions analytically over such a space is beyond the scope of our current work. What we intend to do instead is to provide an approximation of the distribution through discrete sampling of the configurational space. First, we show that if the set of samples fulfill certain conditions, then the estimated distribution approximates the correct distribution.

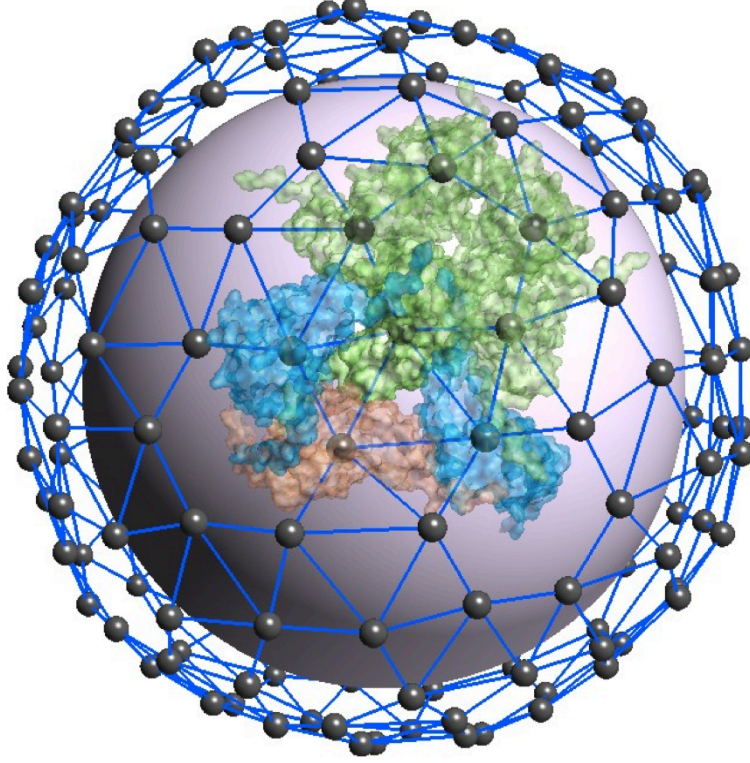


Figure 3: Interface graph for 1CWP (cowpea chlorotic mottle virus). Each node on the graph represents a monomer (some of the monomers are shown). Two nodes are connected if their corresponding monomers share an interface. Any subassembly corresponds to a connected subcomponent of the graph, and the space of local perturbations of the subassembly has $6t$ DOFs, where t is the number of edges in that subcomponent of the graph.

Bounded error of estimation through low-discrepancy sampling For a continuous function f on a d -dimensional product space \mathcal{J}^d , the modulus of continuity is defined as $\omega(f, t) = \sup_{u, v \in \mathcal{J}^d \text{ and } \delta(u, v) \leq t} |f(u) - f(v)|$, where $\delta(u, v)$ is the distance between two configurations/samples. In other words, the value of f does not change without bounds if the parameters are close. Also given a set of N samples $P = \{x_1, x_2, \dots, x_N\}$, we can define their discrepancy with respect a collection of subsets, \mathcal{X} , as:

$$D(P, \mathcal{X}) = \max_{X \in \mathcal{X}} \left(\frac{|P \cap X|}{|P|} - \frac{\mu(X)}{\mu(\mathcal{U})} \right), \quad (1)$$

where μ is the Lebesgue measure (high-dimensional volume), and \mathcal{U} is the universe. Discrepancy can be thought of as the “evenness” of the sample distribution.

Now we present the following theorem:

Theorem 1: Bounded error of integral over \mathcal{J}^d (adapted from Theorem 2.13 of⁴³)

If f is continuous in \mathcal{J}^d , then, for any set of samples $P = \{x_1, x_2, \dots, x_N\}$ such that $x_i \in \mathcal{J}^d$, we have:

$$\left| \int_{\mathcal{J}^d} f(u) du - \frac{1}{N} \sum_{n=1}^N f(x_n) \right| \leq 4\omega\left(f; (D_N^*(P))^{1/d}\right) \quad (2)$$

Essentially, if one ensures that $D_N^*(P)$ is low, then the error of approximation for the integral is bounded. In our case, we want to approximate a distribution. Notice that the above theorem guarantees that if low-discrepancy sampling is performed, the cumulative distribution function (CDF), as well as the moments will be approximated with bounded error.

However, generating such low-discrepancy sampling in a high-dimensional space is nontrivial.

Efficient low-discrepancy sampling in high dimensional spaces Naive approaches in generating low-discrepancy sampling in product spaces involves sampling each degree of freedom uniformly and then combining them in all possible ways to generate samples in high dimensions. This clearly results in exponential number of samples (in terms of d , the dimension of the space), but also does not always guarantee low-discrepancy. Furthermore, it would be computationally intractable to achieve an acceptable level of discrepancy when using such exponential sampling over a space equivalent to \mathbb{R}^{2r+6t} for any practical values of r and t . However, there exist efficient sampling strategies for such spaces which *guarantee* bounded discrepancy using only a polynomial (in the number of degrees of freedom) number of samples.

In this article we leverage the product space sampling algorithm described by Bajaj et al.⁴⁰ which guarantees low-discrepancy defined over subsets of combinatorial rectangles (\mathcal{X} , in the definition in the previous subsection). The basic concept of the algorithm is to use a specific class of pseudo-random number generators (described by Gopalan et al.⁴⁴) that provide strict randomness guarantees. These generators are then used to choose a polynomial-sized subset of samples out of the exponential-sized naive sample space, while still guaranteeing that the discrepancy of the samples is bounded by $d\delta$, d being the number of dimensions and δ being the discrepancy in a

single dimension. It was shown in previous work (see the figure from³⁴) that for dimension greater than about 10, this method of pseudo-random sampling far outperformed traditional methods. As our dimension is $2r + 6t$ (which is much greater than 10 for any practical values of r and t), we have used this method here.

This technique was previously applied by Rasheed et al.³⁴ to bound the uncertainties of different proteins and complexes under large conformational shifts as well as local perturbations. It was found that even with high degrees of freedom, if the range of perturbations and flexible motions are constrained within a neighborhood, a relatively small number of samples are sufficient in providing low approximation error for the distribution of different quantities of interest (e.g. free energy). In this study, we generated 1000 samples for each subassembly.

Computing Distribution of Free Energy and Binding Free Energy

Given a set of samples with low discrepancy, we first compute the free energy for each of the samples. We use Gibbs model of free energy defined as $E = E_{bonded} + E_{vdw} + E_{coul} + G_{cav} + G_{vdw} + G_{pol} - TS$ where E_{bonded} is bonded energy terms representing the energy required to move away from ideal bond lengths, angles, etc., E_{vdw} is the internal van der Waals energy, E_{coul} is electrostatic interaction energy, G_{cav} is approximated using the volume of the protein and the exposed surface area, G_{vdw} is the Van der Waals interaction between exposed atoms and solvent atoms, G_{pol} is the polarization energy (we used Generalized Born approximation), T is the temperature and S is the entropy. In this article, we disregard the effect of S .

We used MolEnergy^{45,46} to compute the molecular surface, area, volume etc. and a fast GPU-accelerated algorithm, PMEOPA⁴⁷ for computing the van der Waals, Coulombic and polarization energies. The accuracy of these techniques were established by Cha et al.⁴⁷ by comparison with AMBER.⁴⁸

While it is trivial to compute binding free energies for static cases simply as the difference of the total free energies before and after binding, it is nontrivial when the input is in the form of distributions, instead. The general idea, however, is still the same, and binding free energies can

be computed as an extension of the single-quantity methodology. First, we define the binding free energies for static cases, as follows: Given a complex or assembly, \mathbf{P} , consisting of a set, \mathbb{S} , of individual chains, we express the binding free energy of \mathbf{P} as:

$$\Delta E(\mathbf{P}) = E(\mathbf{P}) - \sum_{C \in \mathbb{S}} E(C). \quad (3)$$

Now, since each of the components in the above equation is a distribution instead of a scalar, we use a probabilistic definition for the distribution of $\Delta E(\mathbf{P})$. The distribution is approximated based on a collection of 1000 observations. Each observation randomly selects a value from the distribution $E(\mathbf{P})$, and from each distribution $E(C)$ such that $C \in \mathbb{S}$.

Comparing Distributions to Rank Assemblies and Transitions

Analyzing assembly pathways requires the ability to quantify and rank different paths in terms of their likelihood, which is most often related to the binding free energy. Since we are dealing with distributions of such energies (rather than simple scalars), we need a slightly involved technique to compare and rank such distributions. One possible approach is to compare the moments (mean and variance, for instance), but this loses information such as whether the distribution is unimodal or bimodal. The pairwise Wilcoxon signed-rank test⁴¹ uses the entire distribution and provides a way to generate a total ordering among a set of distributions.

For any pair of distributions, X and Y with N points, the Wilcoxon statistic, $W(X, Y)$, is computed as follows:

1. Let $d_i = |x_i - y_i|$, $i \in 1 \dots N$, be the absolute difference between two random, independent draws, $x_i \in X$ and $y_i \in Y$; let σ_i be the sign of $x_i - y_i$
2. Order each d_i from smallest to largest, and let R_i be the rank from this ordering

3. Calculate the Wilcoxon signed-rank statistic, $W(X, Y)$, as:

$$W(X, Y) = \sum_{i=1}^N \sigma_i \cdot R_i \quad (4)$$

It is easy to see that this statistic is symmetric, and that $W(X, Y) \approx -W(Y, X)$.

This can be extended to multiple distributions, X_1, \dots, X_m , as follows:

$$W(X_i | X_1, \dots, X_m) = \sum_{j \neq i}^m W(X_i, X_j) \quad (5)$$

For distributions of energy, a lower value of W is more favorable, thus the most optimal distribution will have the most negative W statistic. Worse (less favorable) distributions will have increasingly positive W statistic.

Results and Discussion

Sufficiency of Sampling

The first and most important question is to determine whether the number of samples we have generated (1000) is sufficient for accurate methods of moments calculations. In previous work,³⁴ we used an incremental sampling approach to determine when our sampling was sufficient enough to obtain confident representative distributions. For all of the single proteins in the study, less than 400 samples was sufficient to provide high confidence in error bounds. As capsid protein complexes consist of many different protein subunits, the uncertainty in a single sample can propagate and influence the stability of the complex as a whole. It is important to ensure we have achieved sufficient sampling.

Table 1 shows the required number of samples across all protein complexes. Most of the subassemblies (86%) required less than 500 samples before saturation was reached. Only 2% of the subassemblies required more than 900 samples, and the most unstable subassembly was

Table 1: Number of samples needed to reach saturation across all samples. Number of samples defined by Chernoff-like bounds for incremental sampling approach. Most of the samples (86%) require less than 500 samples; 79% required less than 400.

# samples needed for saturation	% of subassemblies
<100	2.6
100-200	39.9
200-300	26.0
300-400	10.3
400-500	7.7
500-600	3.3
600-700	3.2
700-800	2.3
900-900	2.2
900-950	1.2
950-975	1.2
>975	0.0

A1-B5-C1-D1-D5, requiring 973 samples. (Upon closer inspection, this subassembly had several outliers that were likely affecting the results.) In our previous study, we showed that even an incremental additive procedure (samples were added until saturation had been reached) with as few as 10 additional samples provided an upper bound on the number of samples needed, so we can safely conclude that the number of samples was sufficient to obtain accurate representative distributions.

Statistical Distribution for Individual Subassemblies

We computed the following quantities for each sample of each subassembly: exposed surface area, enclosed volume, LJ and Coulombic potentials, the solute-solvent polarization energy (G_{pol}), total free energy (GBSA model) and delta energy (ΔG). Sufficiency of the sampling guarantees that the distributions of each of these properties are acceptably accurate. Figure 4 shows the distributions of calculated surface area, exposed volume, energy and G_{pol} for the pentamer, A1-A2-A3-A4-A5. As can be seen in these plots, minor perturbations in internal angles and interface contacts can have large effect on all computed quantities. Some of these changes are intuitive (small changes in internal angles have a large effect on exposed surface area, as seen by the large second moment

of the PDF), but computing the quantities on all samples provides an accurate measurement as to *how much* they can change. Additionally, while most distributions are relatively well-behaved (approximately Gaussian with only one peak), for a small number of subassemblies (especially those with potentially few contacts), the PDF is bimodal, providing additional insight into the stability of the complex (see Supplemental ??????).

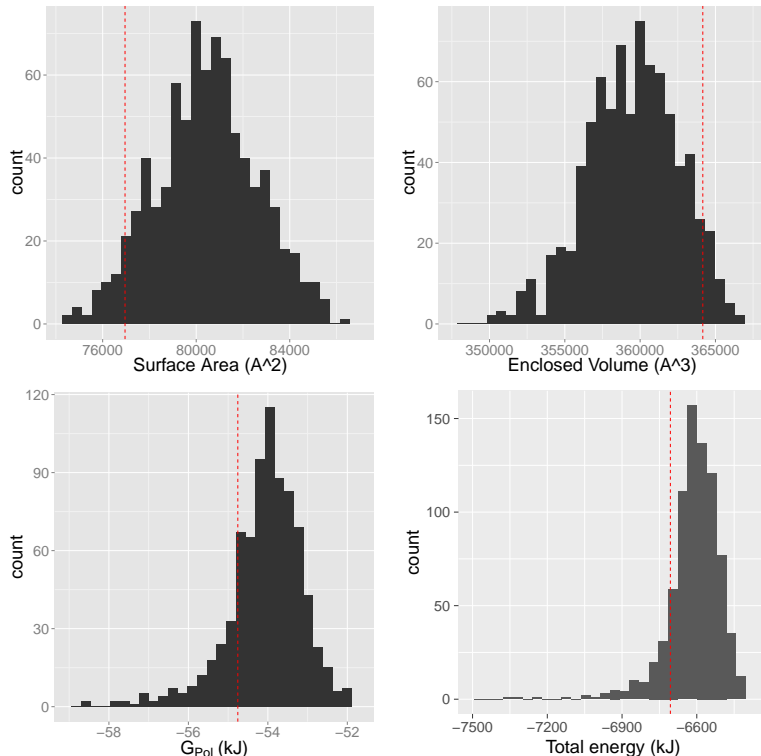


Figure 4: Histogram plots of exposed surface area, enclosed volume, G_{pol} and total energy (GBSA model) for all samples of the pentamer A1-A2-A3-A4-A5. Dotted vertical lines are the quantity computed on the un-sampled molecule.

The second observation from these plots is one major motivation for using distributions of quantities instead of single values. In Figure 4, the dotted red line shows the quantity computed on the non-perturbed subassembly. For the pentamer, the original values vary wildly from the distribution mean. In fact, a histogram of Z-scores for all subassemblies shows that most of them differ greatly from the sampled mean, with many being more than 2 standard deviations away. In addition, subassemblies with a larger number of subunits do not necessarily have a higher variance, so correctly accounting for propagated uncertainty requires good, low-discrepancy samples.

Comparing Capsomers

Given the distributions of a number of subassemblies for a specific property, we can compare them using the Wilcoxon signed rank test and generate a total ordering. This is especially useful in gaining insights (with quantified uncertainty bounds) into the relative binding affinities or stabilities of different capsomers.

Figures 5 and 6 show the distribution of total energy for single subunits and dimer subassemblies, and the top 10 subassemblies of size 3, respectively (see Figure 7 for a surface representation of these complexes). According to this test, the most stable subunit is B, the most stable subassembly of size 2 is the B5-C1 dimer, and the most stable subassembly of size 3 is the A1-B1-B5 complex. The least stable complexes are A, the C1-D1 dimer, and the C1-D1-D7 complex (not pictured).

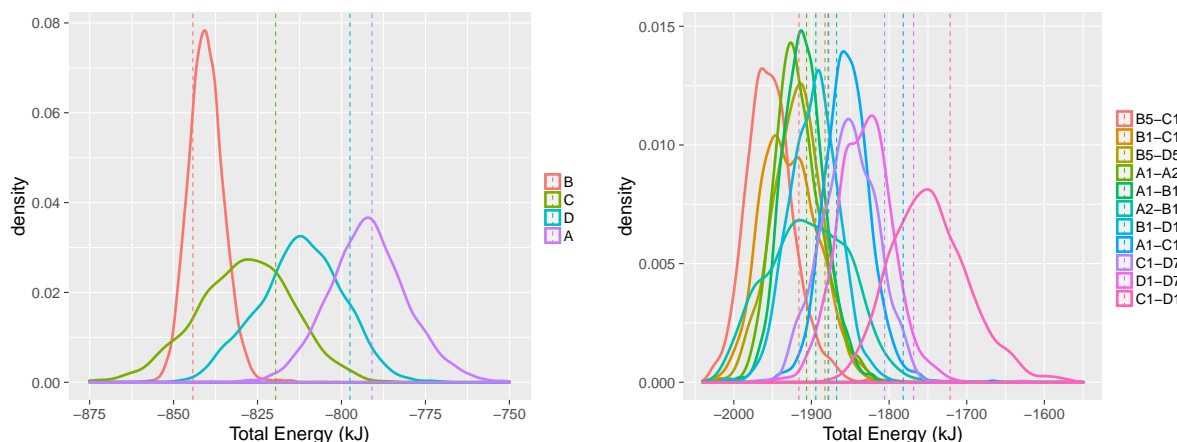


Figure 5: Density plots of distribution of all energy values for all single subunits (left) and subassemblies of size 2 (right). Legend is ranked according to the Wilcoxon signed-rank test (top is best), as reported above.

We also applied the Wilcoxon signed rank test to rank all possible transitions from any given subassembly, which is a crucial step in predicting/analyzing the assembly pathway. For example, Figure 8 shows the possible state transitions starting at C1-B5-B1. If only the native configurations were used, one would have reached the conclusion that adding B7 would be the best transition. However, this is a spurious result since the interface between B7 and B5 have low contact area and is only stable if D7 was also present (see Supplementary Figure S5). That sensitivity is exposed

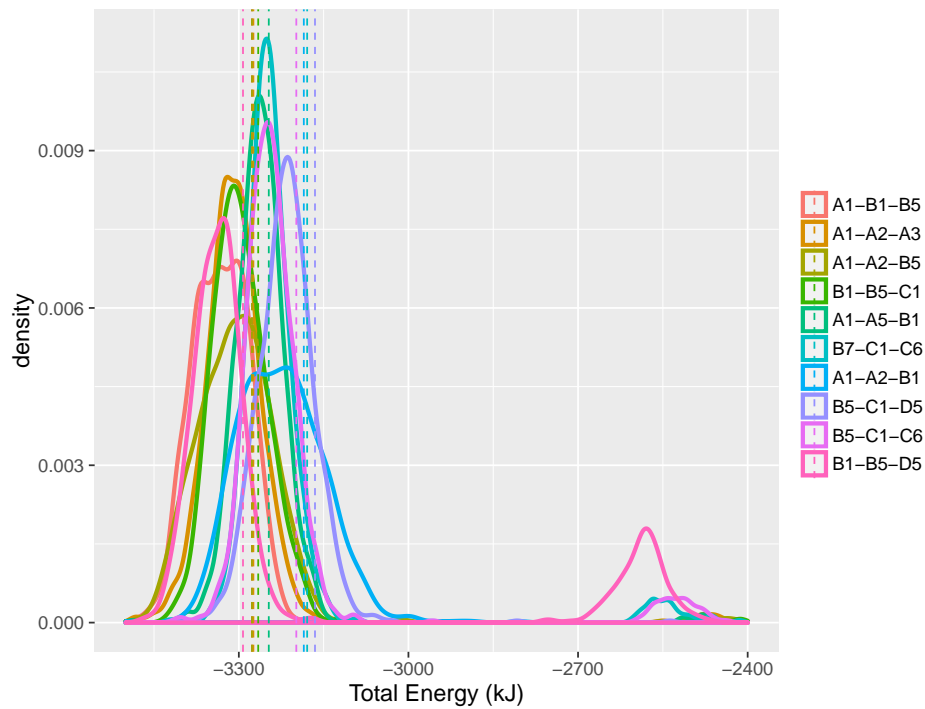


Figure 6: Density plots of distribution of energy values for *top 10* subassemblies of size 3. Legend is ranked according to the Wilcoxon signed-rank test (top is best), as reported above.

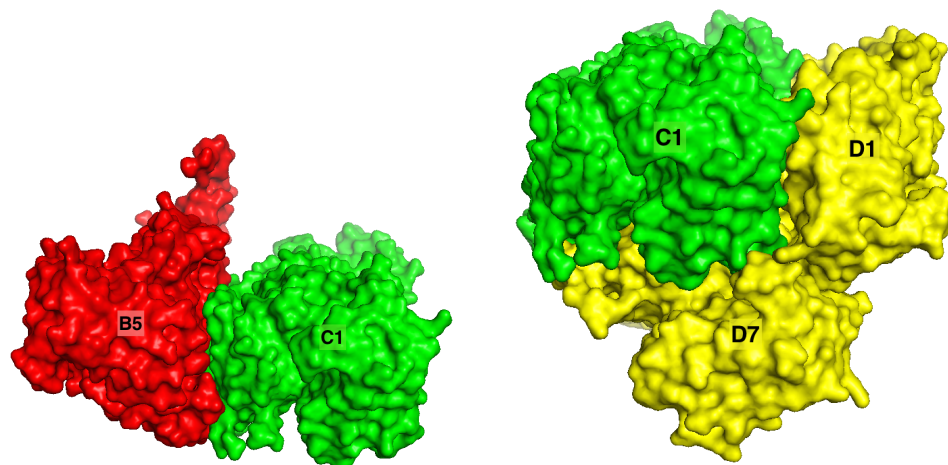


Figure 7: Labeled surface representation of B5-C1 (left), and C1-D1-D7 (right). B5-C1 and C1-D1-D7 are the most stable 2- and 3-subunit capsomers.

through sampling the local configuration space which resulted in some configurations with tighter binding sites and others that pulled them apart (apparent from the bimodal nature of the distribution). Our method successfully accounted for this uncertainty, and the Wilcoxon test is robust enough to determine that it is not the best possible transition.

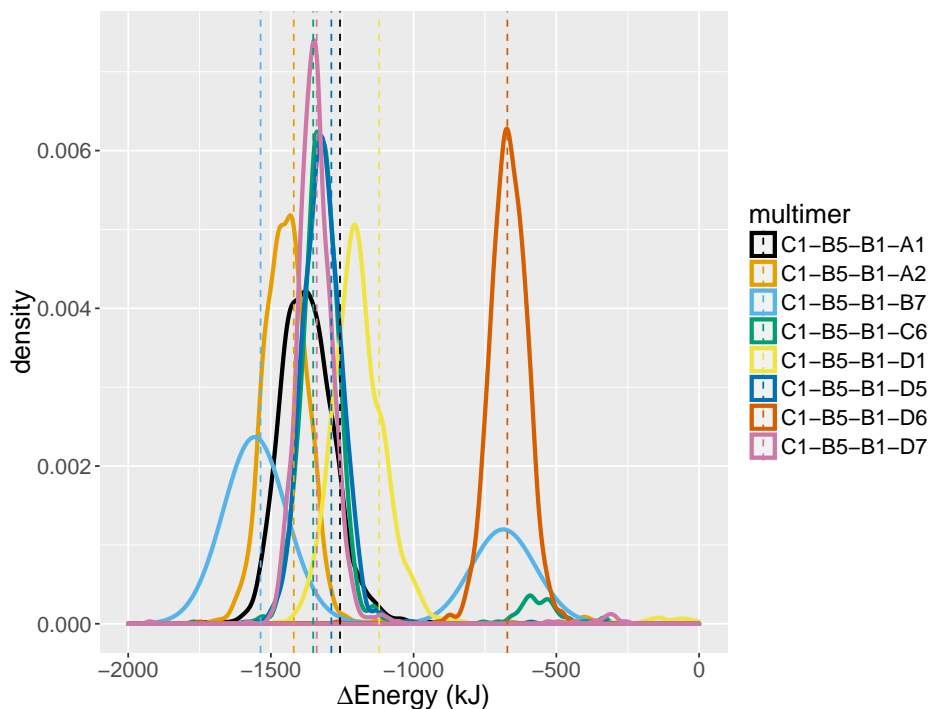


Figure 8: Density plots of distribution of energy values for all possible capsomers starting with C1-B5-B1. Legend is ranked according to the Wilcoxon signed-rank test (top is best), as reported above. Dotted vertical lines show the value computed on only the native structure.

Transitions

Finally, after ranking each transition based off its Wilcoxon score, we construct a complete transition graph showing all possible pathways leading from monomeric subassemblies to the largest subassemblies. Since the entire graph is too large to visually inspect, we present snippets of it here. Figure 9 shows transition sub-graphs starting with subunits B5-C1 and C1-D1-D7 respectively. Note that these two were determined to be very stable states according to our analysis presented in the previous subsection. Starting with (C1-B5), the most likely pathway is (C1-

B5)→B1→A2→A3→A4 (Figure 9, left). For an assembly process starting with subassembly A1-A2, the most likely pathway is (A1-A2)→A3→A4→B5→B1 (Figure 9, right). This kind of figure can provide a visual method for observing likely state transitions. States that are very likely (such as A2-A3-B1-B5-C1, the last star node on the left graph) have many highly-weighted incoming edges (these are "sink" states). States that are not so good can only be reached via poor transitions (such as A1-A5-C1-D1-D7, in the right graph).

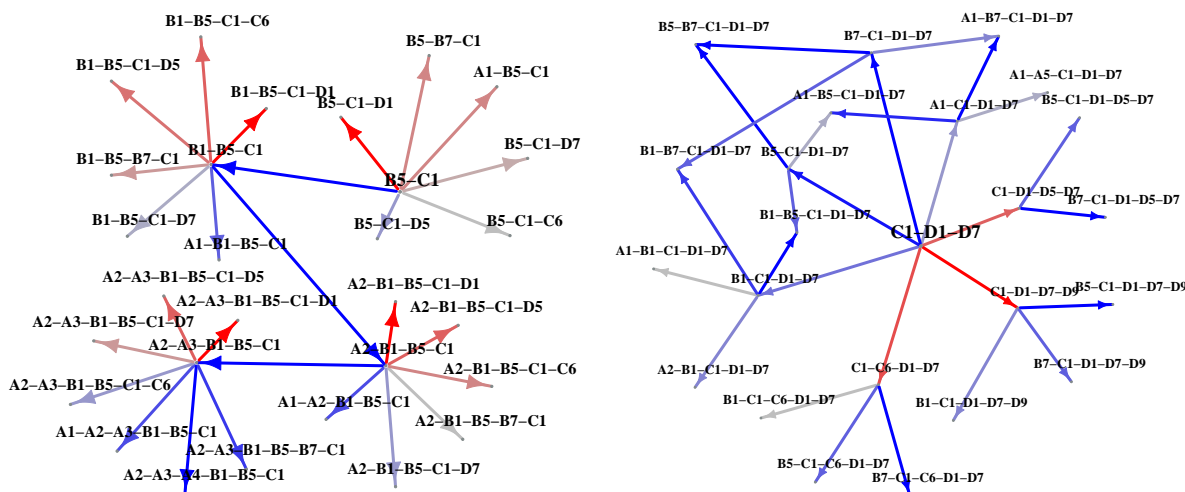


Figure 9: Network graph of possible subassemblies formed when starting from most likely starting points (left: B5-C1 dimer, and right: C1-D1-D7 complex). Color ranges from red (low Wilcoxon score) to grey to blue (high Wilcoxon score). Potential sinks are identified by nodes that have many incoming blue edges, such as A2-B1-B5 on the dimer graph. D1-D7 complex only shows the most likely pathway; C1-D1-D7 graph has low-scoring subassemblies removed from the graph for clarity. For more network graphs of individual subunits, please see Supplemental Figures.

There are several observations that can be made from generating likely pathways. One observation is that, while A1-A2-A3-A4 is a very stable 4-subunit subassembly, A1-A2-A3-A4-B5 is more stable than A1-A2-A3-A4-A5, the pentamer (see Figure 10). This indicates that the pentamer is not fully stabilized until the addition of the B subunit. Such reliance on dimeric interactions may correlate with the size specificity of the virus capsids, since such an interface will not be present in a T=1 shell. Another insight is that the A-C interface is not at all stable (the C-B or C-D interfaces are much better), and probably does not happen until much later in the assembly process or until other partners on hexameric interfaces provide necessary stabilization.

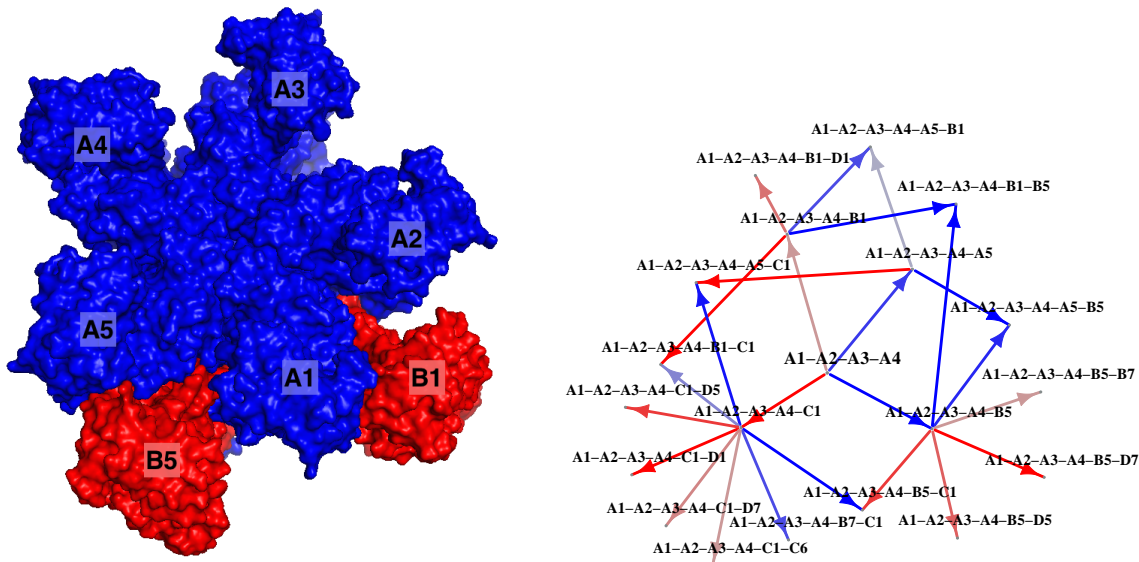


Figure 10: Labeled surface representation of A1-A2-A3-A4-A5-B1-B5 (left), and network graph of possible subassemblies formed when starting from A1-A2-A3-A4. Color ranges from red (low Wilcoxon score) to grey to blue (high Wilcoxon score). According to the Wilcoxon test and our results, the final piece of the pentamer does not form until after B1 and B5 have been added.

Steady-State Concentration Calculations

Transition pathways are useful for determining the most likely step in a single-step reaction. However, single-step reactions do not take into account one of the major driving forces for chemical reactions: concentrations of the necessary reactants and products. If no reactants are available in a chemical reaction, it cannot take place; likewise, if the concentration of products is too high, the reaction will not proceed forward. For this reason, we also sought to provide a *global* view of the viral assembly in terms of concentrations of products and reactants.

According to the Zlotnick,⁴⁹ the rate of formation of a subassembly consisting of n copies of monomer S , S^n , is:

$$\frac{d[S^n]}{dt} = k_{\text{assoc}} [S^{n-1}] [S] - k_{\text{dissoc}} [S^n], \quad (6)$$

where $[S]$ is the concentration of a single subunit, and $[S^{n-1}]$ is the concentration of S^{n-1} without monomer S . At equilibrium ($d[S^n]/dt = 0$), the dissociation expression from Equation ?? indicates

the concentration of free individual subunits that must remain:

$$\frac{[S^n]}{[S^{n-1}][S]} = \frac{k_{\text{assoc}}}{k_{\text{dissoc}}} \quad (7)$$

The ratio $k_{\text{assoc}}/k_{\text{dissoc}}$ at equilibrium is known as the *equilibrium constant*, K_{S^n} , and is directly related to the change in free energy, $\Delta G(S^n)$, as:

$$\Delta G(S^n) = -RT \ln K_{S^n} \quad (8)$$

$$K_{S^n} = \exp\left\{-\frac{\Delta G(S^n)}{RT}\right\}, \quad (9)$$

where R is the gas constant ($8.314 \text{ J mol}^{-1} \text{ K}^{-1}$) and T is the temperature (298K).

Thus, if the concentration of the reactants, $[S]$ and $[S^{n-1}]$, and the change in free energy of the subassembly, $\Delta G(S^n)$, are known, then it is possible to compute the concentration of the product, $[S^n]$:

$$[S^n] = [S^{n-1}][S] * \exp\left\{-\frac{\Delta G(S^n)}{RT}\right\} \quad (10)$$

This can be extended to subassemblies with generic reactants, such as $B5 - C1$ and $A1$, and the product $A1 - B5 - C1$, as long as the concentrations of the reactants and ΔG of the product formation is known.

It should be noted here that for many chemical reactions, the rate of the reaction is determined by kinetics (k_{assoc} and k_{dissoc}) and not by thermodynamics (ΔG). However, when the values of ΔG are high enough, it can be assumed that the reaction will go to completion quickly, and the final ratio of products and reactants is equal to the equilibrium constant. As the ΔG values used in these experiments are very favorable (on the order of 300-3,000kJ/mol), this assumption was made through this section.

Representing Capsid Assembly as a Graph MAP Problem

Based on Equation ??, it is easy to see that the concentration of a single product is dependent on the concentration of one or more reactants. If concentrations of subassemblies are represented by vertices in a graph, then dependencies can be represented by directional edges in this graph, that have weights proportional to the ΔG value for each formation. In this way, we can use a *graphical model* to describe the assembly process. If we initialize this graphical model with non-zero concentrations for the monomers (A, B, C, and D, for 10pM) and zero concentration for all other nodes and allow concentration to “flow” along edges from one node to another, then the *maximum a posteriori probability (MAP) estimate* is the steady-state of the graph, where no flow is happening.

As the details of this representation are beyond the scope of this work, we will only state that we are using a Bayesian factor graph to represent the steady-state assembly and leave further details for the interested reader to the supplement of this paper. For figures in this paper, the Nudaurelia Capensis virus capsid was modeled with initial concentration of all monomers (A, B, C, and D) were typical concentrations at micromolar ranges (100e-nM). The message-passing algorithm was run until concentration change was below 1e-11M, summed over all concentration nodes.

Steady-State Concentrations

Figure 11 shows the distribution over successive steps of the message passing algorithm, which attempts to model the self-assembly of the capsid. For the most part, the assemblies with higher concentrations at the steady state (final step of the algorithm) are those with more subunits (e.g., the concentration of the hexamer B5-B7-C1-C6-D5-D7 was 26nM, several orders of magnitude higher than the other subassemblies). This observation would suggest that the subassembly formation largely proceeded toward completion of products, and that the limiting factor was concentration of the products, as was expected.

An important point to note is that this graph would be different if subassemblies of size greater than 6 were also included, as the intermediate products would be quickly consumed. From Fig-

ure 11, this phenomenon can already be observed; for example, the intermediate product A1-C1 has an initial high concentration, but then quickly drops off as it is used for later products, such as A1-B1-C1-D1. This also explains why the concentration of monomer C decays so slowly, as there are fewer beneficial reactions involving C (see, for e.g. Figure 6, where the products involving C are not highly ranked). This might suggest that the configuration of C with the rest of the capsid is meant as a stabilizing subassembly, and is not used until much later in the assembly process.

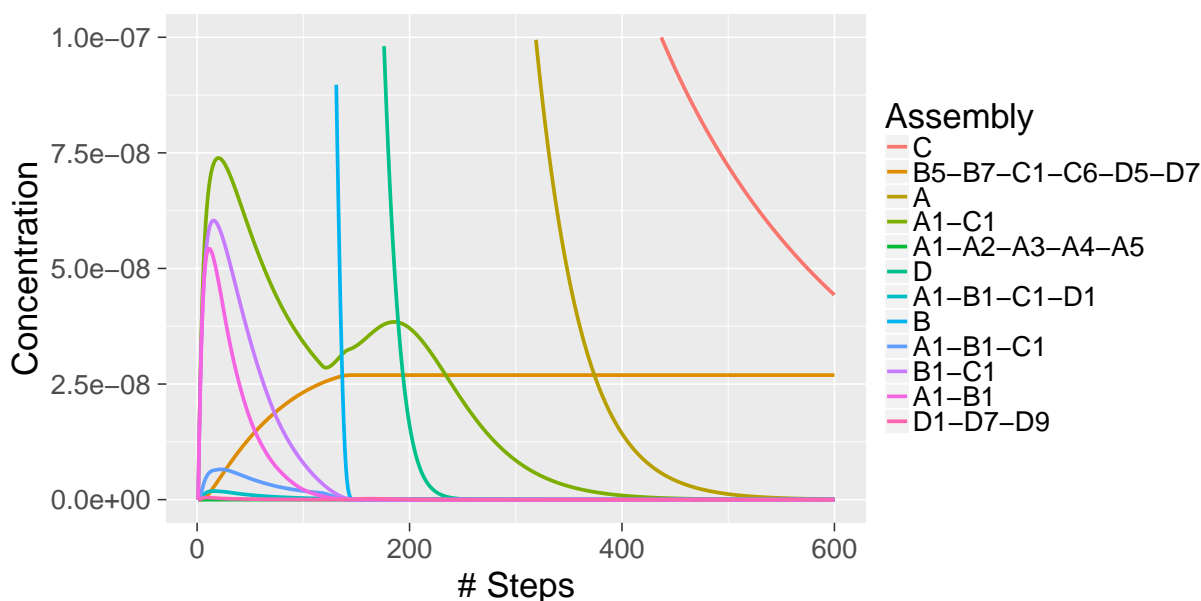


Figure 11: Change in concentration over time for several subassemblies of the Nudaurelia Capensis viral capsid. Legend is ranked according to concentration at step 600. The x-axis has been trimmed to emphasize initial concentration changes, as the steady-state concentrations were reached after 1300 steps of the algorithm.

Finally, we can also plot the distribution of all possible steady state concentrations, shown in Figure 12. For this plot, initial values of ΔG were taken from the *distribution* of possible values for each subassembly, and then the steady-state assembly algorithm was run as usual. Distribution of final concentration of subassemblies were plotted as box-and-whisker plots. The red dot shows the value computed when using ΔG computed on the unperturbed PDB structure. This plot shows several things: first, that the distributions of final concentrations differ greatly across subassemblies, and second, that the value computed on the original PDB does not represent the true average across all samples.

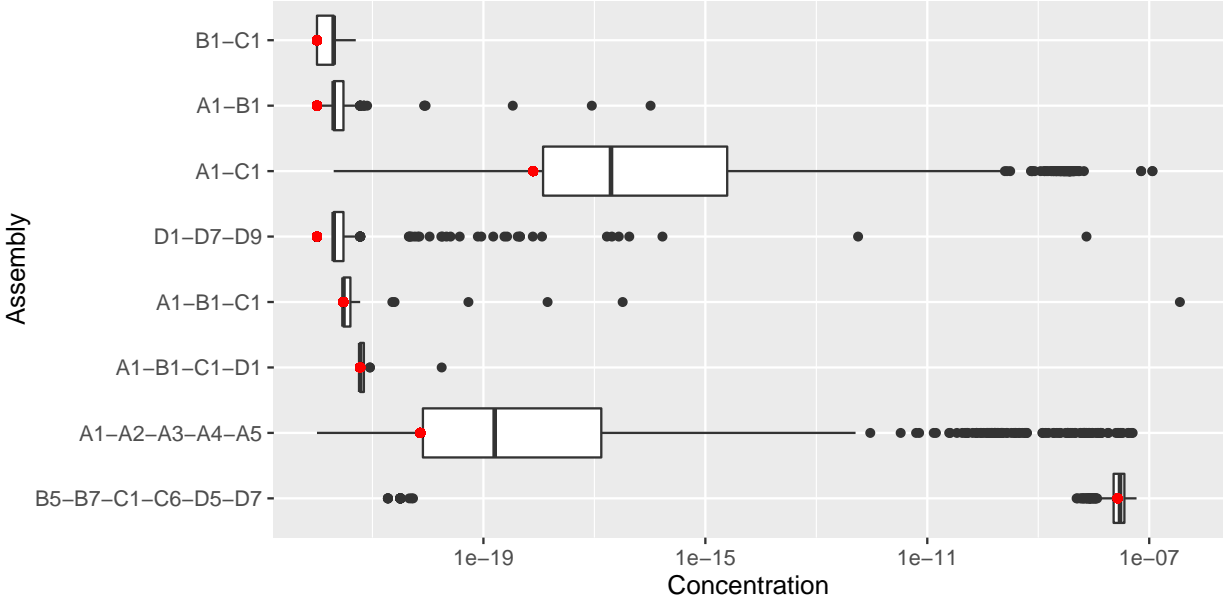


Figure 12: Distribution of concentrations over different input ΔG values, for various subassemblies of the *Nudaurelia Capensis* viral capsid. X-axis is log-scale to emphasize differences. Red point is value computed when using just the original PDB.

Conclusions

Most of existing research in assembly pathway prediction/analysis of virus capsids have relied on the final configuration of the capsid to determine the configuration of the intermediate states. This assumption is overly simplified since the capsid proteins may undergo conformational changes, binding interfaces may change to allow binding with another subassembly, etc. throughout the assembly process until stabilization. To better capture this phenomenon, we have developed a statistical-ensemble based approach which sufficiently samples the configurational space of each monomer and the relative local orientation between monomers to capture the uncertainties in their binding. Essentially, instead of modeling each subassembly as a static configuration, they are modeled as distributions of possible configurations. This allows us to compute the free energy of a subassembly, and the binding free energy on a possible assembly edge as distributions so that statistical guarantees of accuracy can be additionally derived for each of the resulting assemblies.

Unlike traditional approaches where pentamers+hexamers, or trimers are used as fundamental building blocks in the assembly pathway analysis, we use individual monomers as our starting

constituents, and consider all possible unique subassemblies (modulo symmetry), of sizes up to 6. The primary aim is to quantitatively understand the formation of the larger building blocks (i.e. trimers, pentamers and hexamers).

We additionally, adapted the Wilcoxon measure to provide a way to compare the distributions and determine the most likely subassemblies that can be generated in any step. Using this score as weights on an assembly graph revealed that there are some low-energy subassemblies that are unlikely to be formed because there are poor transitions along the path that forms said subassembly. Finally, we extended the assembly prediction to factor concentrations of the constituents. The assembly graph was converted to a Bayesian factor graph where the final concentrations of the subassemblies are posed as a graphical maximum a posteriori problem. Transition probabilities were set up based on the equilibrium constant computed from the binding free energy, and both forward (association), and backward (dissociation) reactions were allowed. The result showed expected patterns, e.g. dimers A1-B1, A1-A2 etc. getting produced at a fast rate and then being consumed as other subassemblies become available, forming the larger subassemblies A1-B1-C1 (trimer), A1-A2-A3-A4-A5 (pentamer), etc. As the concentrations reach their steady state, larger particles had higher final concentrations, as was expected.

In summary, we contend that the use of *ensemble distributions of molecules*, instead of single conformations, allows one to make statistical inferences about the stability of molecular subassemblies. We have shown that, a full distribution of possible subassemblies is not obtainable if one was to use assembly combinations only from the original PDB conformation. This could often lead to erroneous conclusions. Use of a statistically rigorous procedure such as the one advocated in this paper, yields inferences on capsid assembly can be made with confidence.

Supporting Information Available

We have added several files containing more in depth description of our methods, additional figures, and some data as supporting information. A list of the files and there content follows:

- Supporting Information for Publication: Contains details of the Bayesian Factor Graph construction, and additional figures S1-S10.
- `graphical_model_edges.avg.txt`: The graphical model for MAP analysis in ASCII format.
- `Wilcoxon_graph.txt`: The complete graph corresponding to Figure 9.

Acknowledgements

This research was supported in part by NIH grant R01-GM117594-0 and Sandia subcontract SNL-1439100.

Supplemental Information

Representing Viral Assembly as Bayesian Factor Graph

For a simple virus ($T=1$) consisting of n identical subunits, it is easy to represent the formation of a virus as a Bayesian network consisting of m nodes, where each node represents a possible subassembly, s_k , and the transition probabilities are the rates of formation. A node, v_{s_k} , representing s_k , would have incoming edges from all s_i and s_j , where $i + j = k$ (e.g. s_{k-1} and s_1 , as well as s_{k-2} and s_2 , etc.). See Figure S1 (left), where a larger subassembly, s_4 , is created from the combination of $s_1 + s_3$ or $s_2 + s_2$. Edges leading into s_4 are labeled with the specific combination. The concentration $[s_4]$ is dependent on $[s_3]$, $[s_2]$, and $[s_1]$, as well as the corresponding ΔG values.

Bayesian Factor Graph One of the limitations with a traditional Bayesian network is that a simple edge weight does not contain all the information necessary to determine forward and backward formation. The dependencies of s_4 , for example, are s_1 , s_2 , and s_3 . However, s_4 is constructed of two s_2 subunits, but only one s_1 subunit. Instead of a Bayesian network, we can represent the viral

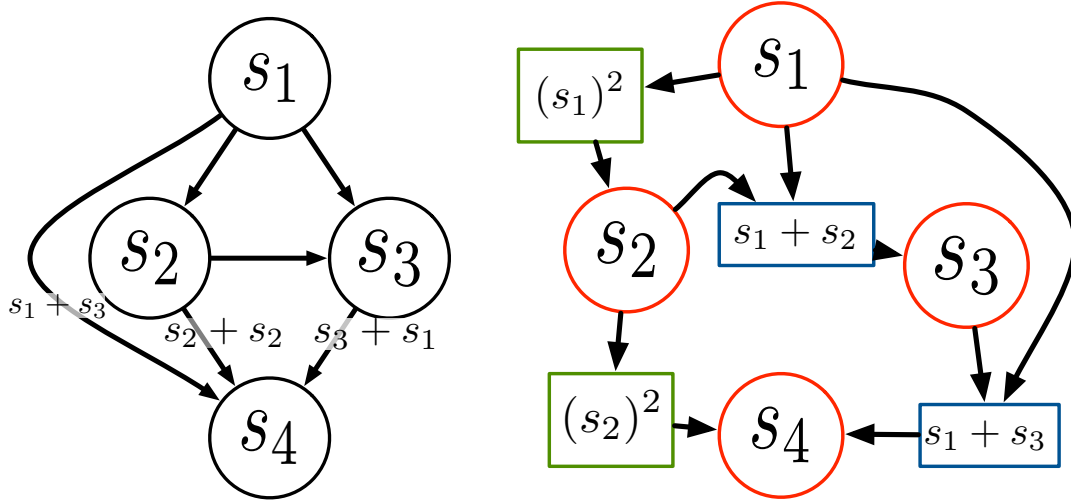


Figure S1: Two different graphical representations of virus formation. Left: a simple directed Bayesian network; right: the same network represented by a factor graph, where *factor* nodes are rectangles, representing the combination of a set of smaller nodes.

assembly process as a Bayesian factor graph, where the *factor* nodes are states that contain this additional information. Thus, the only dependencies for s_4 are $s_2 + s_2$ or $s_1 + s_3$, not the individual subunits. Because we are modeling the creation of a virus from the addition of single subunits, we will enforce an additional constraint that each factor node must have exactly two incoming edges and exactly one outgoing edge.

Solving MRFs Using Message Passing

The traditional method for solving MRFs is through a technique called *belief propagation* or *message passing*, where each node in the graph with “propagate” its “belief” about the current state of the network to its neighbors.⁵⁰ Traditional belief propagation is performed by the sum-product method,⁵¹ which will be adapted for our use here.

Let r be a subunit that is used to produce more than one product, e.g. r and r_1 form p_1 , and r and r_2 form p_2 . Then the ratio of the two products can be determined from Equation ???. If we

assume that the concentration of r_1 and r_2 are equal, then:

$$\frac{[p_1]}{[p_2]} = \frac{\exp\left\{-\frac{\Delta G(p_1)}{RT}\right\}}{\exp\left\{-\frac{\Delta G(p_2)}{RT}\right\}} \quad (11)$$

$$= \exp\left\{\frac{1}{RT} (\Delta G(p_2) - \Delta G(p_1))\right\} \quad (12)$$

If we let $e(p_2)$ be the amount in the exponent, i.e.:

$$e(p_2) = \exp\left\{-\frac{\Delta G(p_2)}{RT}\right\}, \quad (13)$$

then we can rewrite Equation ?? as:

$$\frac{[p_1]}{[p_2]} = \frac{e(p_1)}{e(p_2)} \quad (14)$$

For a set of k potential products, $p_1 \dots p_k$, the proportional concentration of reactant r that will be used to form product p_i , $\lambda(p_i)$, can be written from Equation ?? and ?? as:

$$\lambda(p_i) = \frac{e(p_i)}{\sum_{j=1}^k e(p_j)} \quad (15)$$

For the sake of notation, we will also define the reverse exponent amount for reactant r , $e^{-1}(r)$ as

$$e^{-1}(r) = \exp\left\{\frac{\Delta G(r)}{RT}\right\} \quad (16)$$

and the reverse proportion of a reactant, $\lambda^{-1}(r_j)$, over a set of potential reactants, $r_1 \dots r_k$, as:

$$\lambda^{-1}(r_j) = \frac{e^{-1}(r_j)}{\sum_{i=1}^k e^{-1}(r_i)} \quad (17)$$

Let $G = [E, V, F]$ be a bipartite graph, G , with edges E and nodes divided into two disjoint groups V (variable nodes, or concentrations of subassemblies) and F (factor “pool” nodes, or hidden nodes where concentrations ‘pool’ before being assembled). Then the only messages that are passed are from $v \in V$ to $a \in F$ and vice versa, and not between two members of the same set.

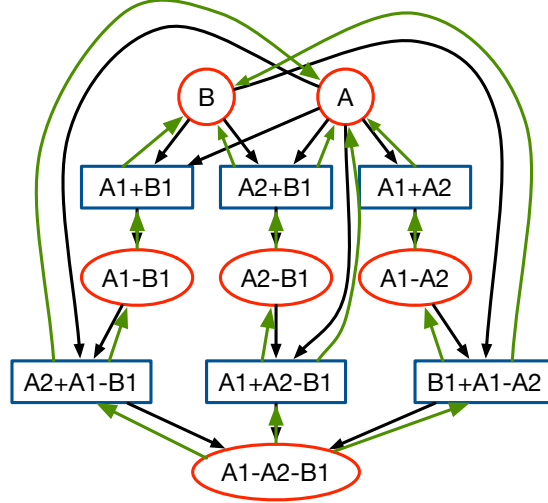


Figure S2: Factor graph for the formation of $A1 - A2 - B1$ from all possible units. Variable nodes are red circles, factor nodes are blue rectangles, forward edges are black, and backward edges are green.

We will also define two different kinds of edges, *forward*, and *backward*. Forward edges represent the formation of a product by two reactants, and backward edges the break-down of products into reactants. See Figure S2 for an example of the formation of $A1 - A2 - B1$ from all possible reactants. The we will define four types of messages here:

Messages along the forward edges are:

- $\mu_{v \rightarrow a}([v])$: A message along from a variable (concentration) node to a factor (pool) node describing the belief of $[v]$, that will be available at pool node a . Pool node a will be used to pool the subunits for subassembly $v + x^*$, where x^* is some (possibly different) subunit.

$$\mu_{v \rightarrow a}([v]) = \lambda(v - x^*) * [v], \quad (18)$$

where $\lambda([v - x^*])$ is defined according to Equation ??, the predicted amount of v that will be available at pool node a .

- $\mu_{a \rightarrow v}([v])$: A message from factor (pool) node a to variable (concentration) node v will be the amount of v that is created. Since each factor node has exactly two incoming edges, let

us call these edges $a[1]$ and $a[2]$. Then the message passed to v will be:

$$\mu_{a \rightarrow v}([v]) = a[1]a[2] * e(v) \quad (19)$$

Messages along the backward edges are:

- $\mu_{v_{a[j]} \leftarrow a}([v_{a[j]}])$: Because the forward message from a reactants pool node to the products concentration node will not always consume all of the products, it becomes necessary to send a message back to the reactants containing the concentration amount not used. This message will be:

$$\mu_{v_{a[j]} \leftarrow a}([v_{a[j]}]) = a[j] - \min(a[1], a[2]), \quad j = \{1, 2\} \quad (20)$$

- $\mu_{a_i \leftarrow v}([a_i])$: The concentration at product v will come from several sources; the total product must therefore reach equilibrium with all of them. Thus, the message back from the product node v to the factor node a about product $a[1]$ or $a[2]$ will be:

$$\mu_{a[1] \leftarrow v}([a_1]) = \mu_{a[2] \leftarrow v}([a_2]) = \lambda^{-1}(v) * [v], \quad (21)$$

where $\lambda^{-1}(v) * [v]$ is defined according to Equation ??.

Message Passing Algorithm The message passing algorithm will be run in the following steps. Let T_0 be the initial concentration of monomer subunits, and let k be the size (number of monomers) of the largest subassembly.

1. Initialize monomers to starting concentration, T_0
2. Until no changes in concentrations have been made:
 - (a) For each size $s = 1$ to $s = 6$ (monomer to largest):
 - i. pass messages along forward edges from v of size s to f of size $s + 1$

- ii. pass messages along forward edges from f of size $s + 1$ to v of size $s + 1$
- (b) For each variable node, v_i , set the concentration $[v_i]$ to the sum of all messages along forward edges, $\mu_{* \rightarrow v_i}([v_i])$
- (c) For each size $s = 6$ to $s = 1$ (largest to monomer):
 - i. perform message along all *backward* edges from v of size s to f of size s
 - ii. perform message passing along *backward* edges from all f of size s to v of size $s - 1$
- (d) For each variable node, v_i , add $\mu_{a \rightarrow v_i}^*([v_i])$ to the current concentration, $[v_i]$
- 3. Report the concentration of all v nodes

Supplemental Figures

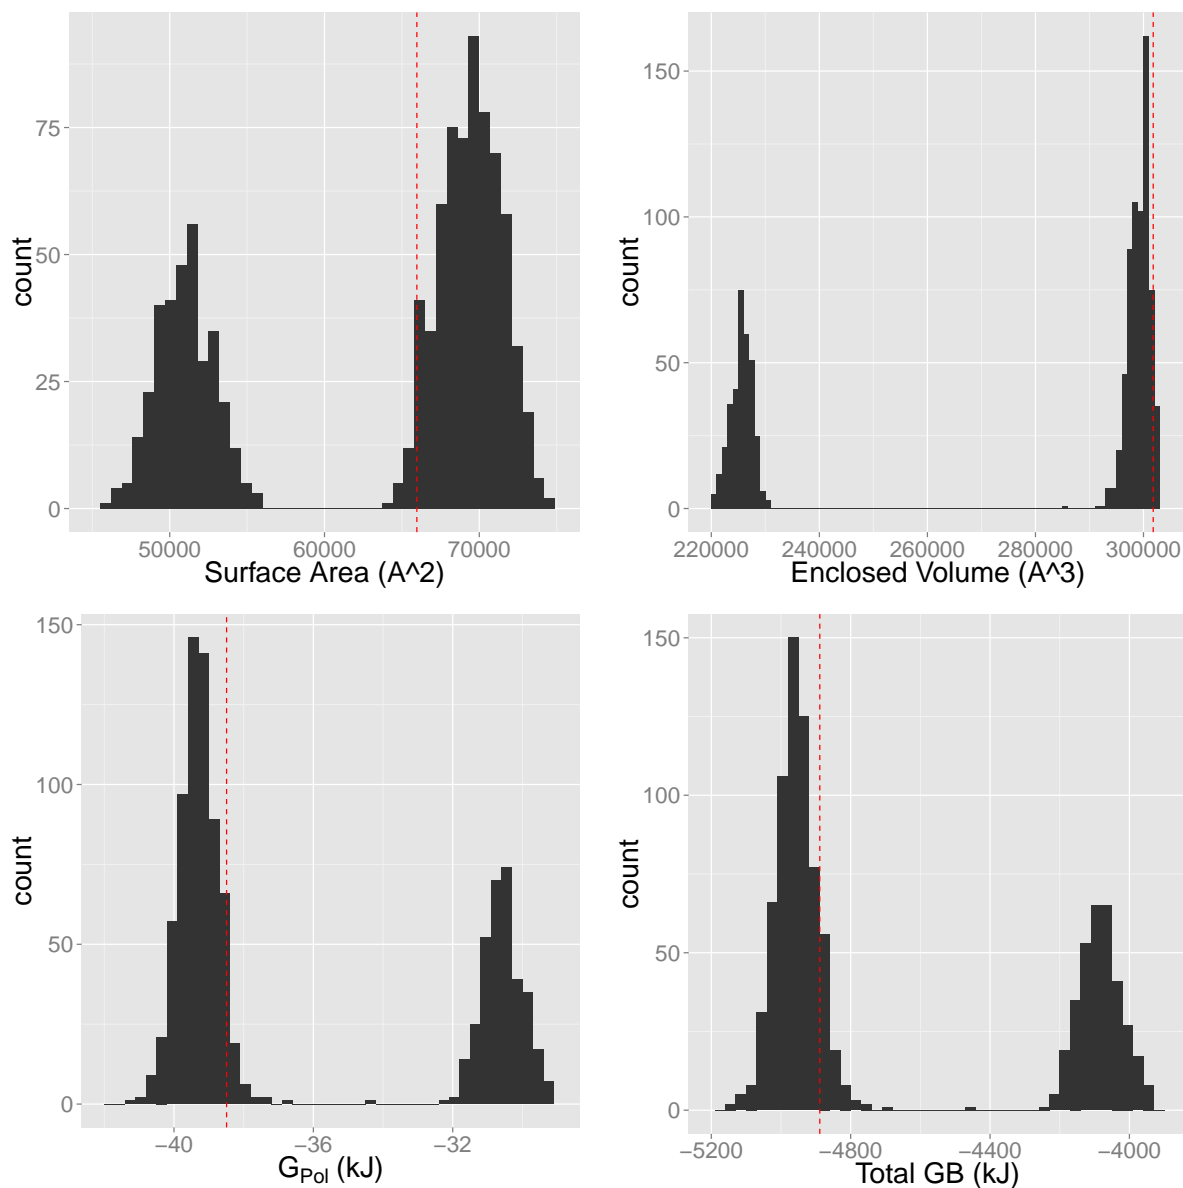


Figure S3: Histogram plots of exposed surface area, enclosed volume, G_{pol} and total energy (GBSA model) for all samples of B1-B5-B7-C1. Dotted vertical lines show the value computed on the un-sampled molecule. Bimodal nature of plots reveal low surface area between B7 and the rest of the complex. See the surface representation in S5.

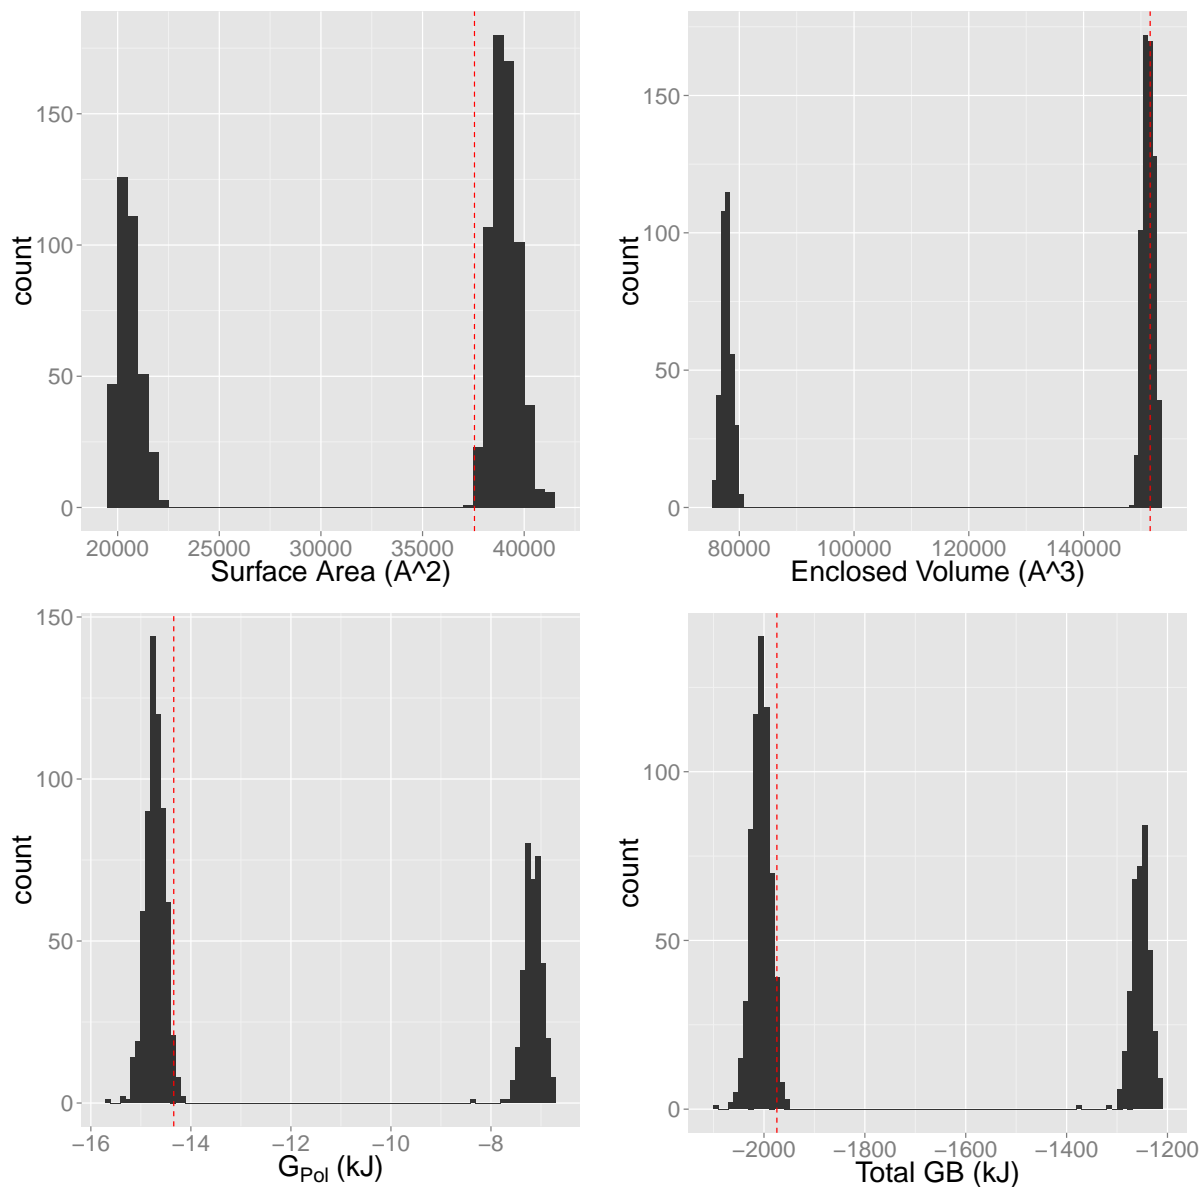


Figure S4: Histogram plots of exposed surface area, enclosed volume, G_{pol} and total energy (GBSA model) for all samples of B5-D7. Dotted vertical lines show the value computed on the un-sampled molecule. Bimodal nature of plots reveal low surface area between B5 and D7. See the surface representation in Figure S5.

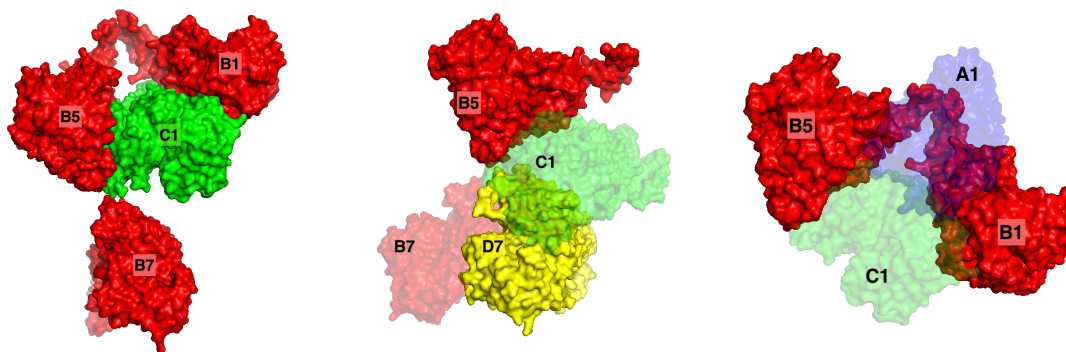


Figure S5: Labeled surface representation of B1-B5-B7-C1 (left), B5-D7 (center), and B1-B5 (right). Complexes have been rotated to show interfaces.

References

- (1) Lai, Y.; K., T.; Sawaya, M.; Asturias, F.; Yeates, T. Structure and Flexibility of Nanoscale Protein Cages Designed by Symmetric Self-Assembly. *J. Am. Chem. Soc.* **2013**, *135*, 7738–7743.
- (2) Zochowska, M.; Paca, A.; Schoehn, G.; Andrieu, J.; Chroboczek, J.; Dublet, B.; E., S. Adenovirus Dodecahedron, as a Drug Delivery Vector. *PLoS One* **2009**, *4*, 1–12.
- (3) Cheng, Y.; Samia, A.; Meyers, J.; Panagopoulos, I.; Fei, B.; Burda, C. Highly Efficient Drug Delivery with Gold Nanoparticle Vectors for in Vivo Photodynamic Therapy of Cancer. *J. Am. Chem. Soc.* **2008**, *130*, 10643–10647.
- (4) Smith, M. T.; Hawes, A. K.; Bundy, B. C. Reengineering Viruses and Virus-like Particles Through Chemical Functionalization Strategies. *Curr. Opin. Biotechnol.* **2013**, *24*, 620–626.
- (5) Helgstrand, C.; Munshi, S.; Johnson, J. E.; Liljas, L. The Refined Structure of Nudaurelia Capensis ω Virus Reveals Control Elements for a $T = 4$ Capsid Maturation. *Virology* **2004**, *318*, 192–203.
- (6) Hagan, M. F.; Chandler, D. Dynamic Pathways for Viral Capsid Assembly. *Biophys. J.* **2006**, *91*, 42–54.

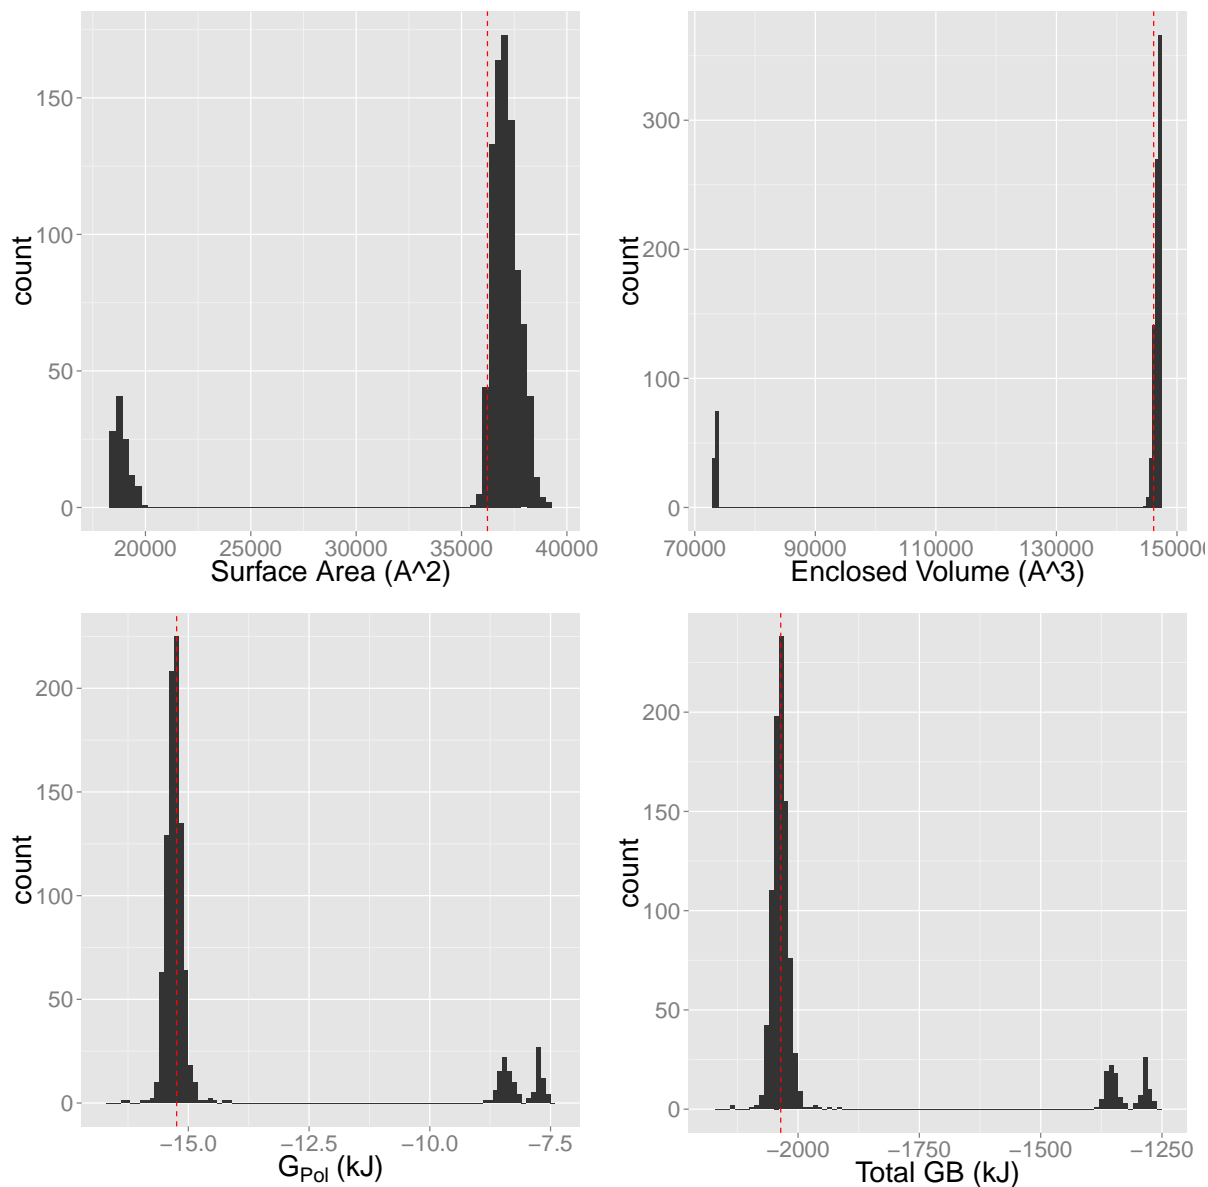


Figure S6: Histogram plots of exposed surface area, enclosed volume, G_{pol} and total energy (GBSA model) for all samples of B1-B5. Dotted vertical lines show the value computed on the un-sampled molecule. Bimodal nature of plots reveal low surface area between B1 and B5. See the surface representation in Figure S5.

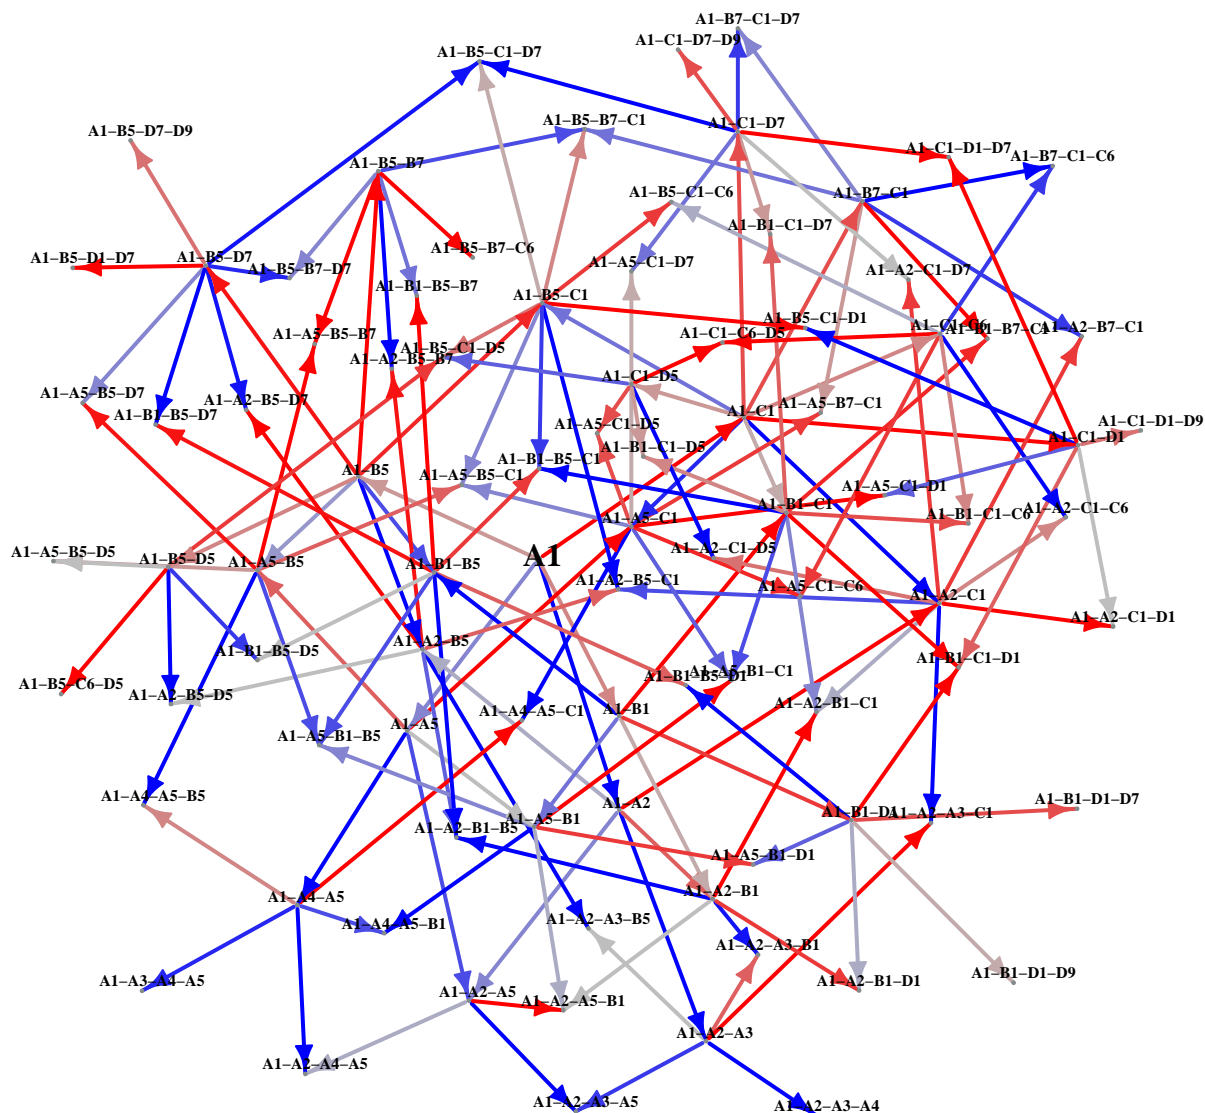


Figure S7: Network graph of possible subassemblies formed, up to size 4, when starting from subunit A1. Color ranges from red (low Wilcoxon score) to grey to blue (high Wilcoxon score). Potential sinks are identified by nodes that have many incoming blue edges, and nodes with low probability of occurring have many incoming red edges.

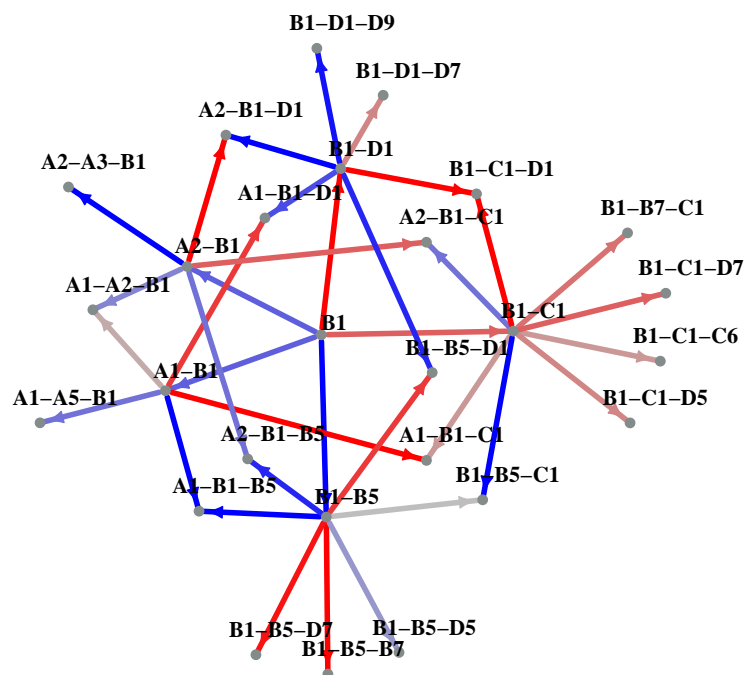


Figure S8: Network graph of possible subassemblies formed, up to size 3, when starting from subunit B1. Color ranges from red (low Wilcoxon score) to grey to blue (high Wilcoxon score). Potential sinks are identified by nodes that have many incoming blue edges, and nodes with low probability of occurring have many incoming red edges.

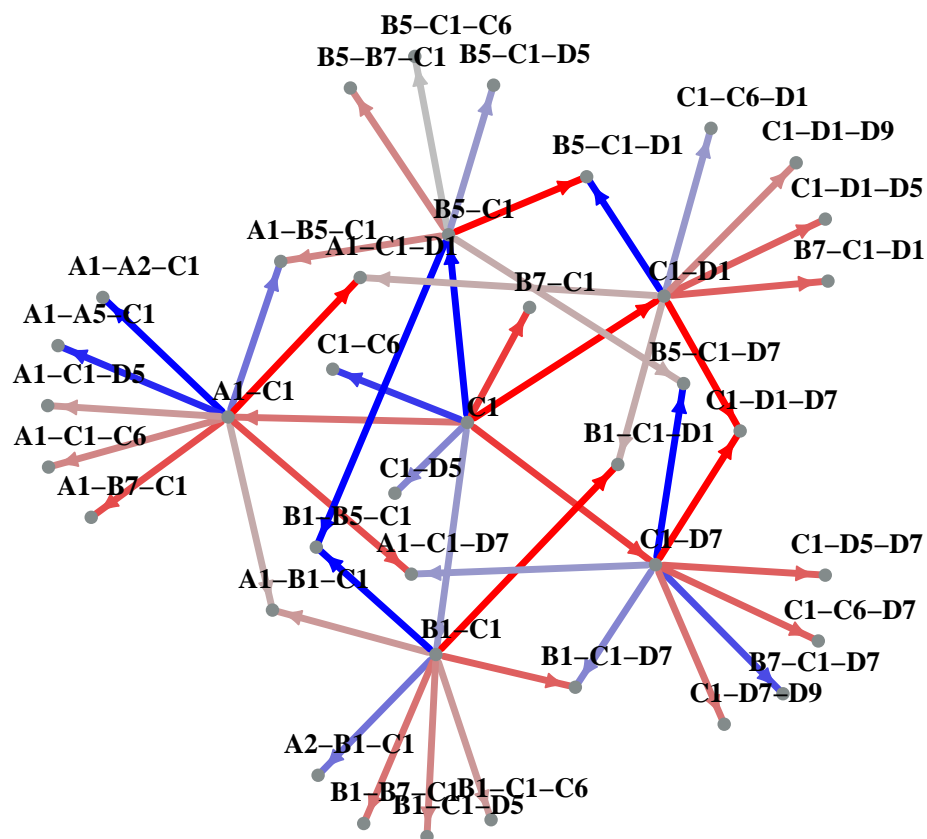


Figure S9: Network graph of possible subassemblies formed, up to size 3, when starting from subunit C1. Color ranges from red (low Wilcoxon score) to grey to blue (high Wilcoxon score). Potential sinks are identified by nodes that have many incoming blue edges, and nodes with low probability of occurring have many incoming red edges.

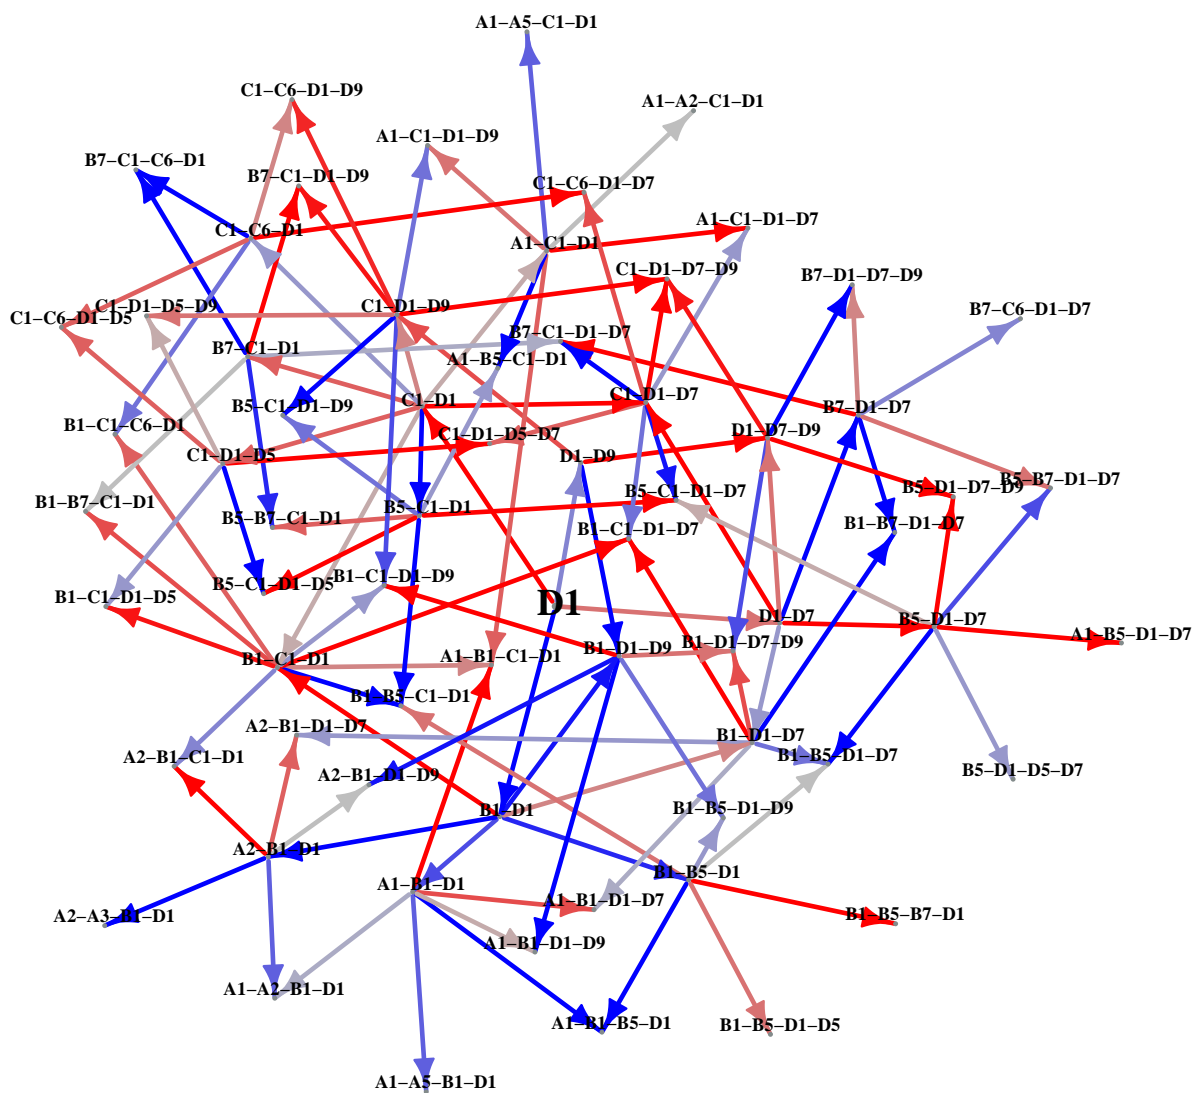


Figure S10: Network graph of possible subassemblies formed, up to size 4, when starting from subunit D1. Color ranges from red (low Wilcoxon score) to grey to blue (high Wilcoxon score). Potential sinks are identified by nodes that have many incoming blue edges, and nodes with low probability of occurring have many incoming red edges.

- (7) Hagan, M. F. Modeling Viral Capsid Assembly. *Adv. Chem. Phys.* **2014**, *155*, 1.
- (8) Zlotnick, A. To Build a Virus Capsid: An Equilibrium Model of the Self Assembly of Polyhedral Protein Complexes. *J. Mol. Biol.* **1994**, *241*, 59–67.
- (9) Zlotnick, A.; Mukhopadhyay, S. Virus Assembly, Allostery and Antivirals. *Trends Microbiol.* **2011**, *19*, 14–23.
- (10) Rapaport, D. C.; Johnson, J. E.; Skolnick, J. Supramolecular Self-assembly: Molecular Dynamics Modeling of Polyhedral Shell Formation. *Comput. Phys. Commun.* **1999**, *121-122*, 231–235.
- (11) Caspar, D. L.; Klug, A. Physical Principles in the Construction of Regular Viruses. Cold Spring Harbor Symposium on Quantitative Biology. 1962; pp 1–24.
- (12) Goldberg, M. A Class of Multi-Symmetric Polyhedra. *Tohoku Mathematical Journal* **1937**, 104–108.
- (13) Janner, A. Towards a Classification of Icosahedral Viruses in Terms of Indexed Polyhedra. *Acta Crystallogr.* **2006**, *A62*, 319–330.
- (14) Keef, T.; Twarock, R. Affine Extensions of the Icosahedral Group With Applications to the Three-dimensional Organisation of Simple Viruses. *Journal of Mathematical Biology* **2009**, *3*, 287–313.
- (15) Pawley, G. Plane Groups on Polyhedra. *Acta Crystallogr.* **1961**, *15*, 49–53.
- (16) Rasheed, M.; Bajaj, C. Highly Symmetric and Congruently Tiled Meshes for Shells and Domes. *Procedia Eng.* **2015**, *124*, 213–225.
- (17) Damodaran, K.; Reddy, V. S.; Johnson, J. E.; Brooks, C. L. A General Method to Quantify Quasi-equivalence in Icosahedral Viruses. *J. Mol. Biol.* **2002**, *324*, 723–737.

- (18) Carrillo-Tripp, M.; Brooks, C. L.; Reddy, V. S. A Novel Method to Map and Compare Protein-protein Interactions in Spherical Viral Capsids. *Proteins* **2008**, *73*, 644–655.
- (19) Mannige, R. V.; Brooks, C. L.; III, T. T. T. Tilable Nature of Virus Capsids and the Role of Topological Constraints in Natural Capsid Design. *Phys. Rev. E: Stat., Nonlinear, Soft Matter Phys.* **2008**, *77*, 051902–9.
- (20) Mannige, R. V.; Brooks, C. L. Periodic Table of Virus Capsids: Implications for Natural Selection and Design. *PLoS One* **2010**, *5*, 7.
- (21) Rapaport, D. Self-assembly of Polyhedral Shells: A Molecular Dynamics Study. *Phys. Rev. E* **2004**, *70*, 1–13.
- (22) Rapaport, D. C. Modeling Capsid Self-assembly: Design and Analysis. *Phys. Biol.* **2010**, *7*, 045001.
- (23) Elrad, O. M.; Hagan, M. F. Mechanisms of Size Control and Polymorphism in Viral Capsid Assembly. *Nano Lett.* **2008**, *8*, 3850–3857.
- (24) Berger, B.; Shor, P. W.; Tucker-Kellogg, L.; King, J. Local Rule-Based Theory of Virus Shell Assembly. *Proc. Natl. Acad. Sci. U. S. A.* **1994**, *91*, 7732–7736.
- (25) Berger, B.; Shor, P. W. *On the Mathematics of Virus Shell Assembly*; 1994.
- (26) Xie, L.; Smith, G. R.; Feng, X.; Schwartz, R. Surveying Capsid Assembly Pathways Through Simulation-based Data Fitting. *Biophys. J.* **2012**, *103*, 1545–1554.
- (27) Schwartz, R.; Shor, P.; Prevelige, P.; Berger, B. Local Rules Simulation of the Kinetics of Virus Capsid Self-assembly. *Biophys. J.* **1998**, *75*, 2626–2636.
- (28) Bona, M.; Sitharam, M.; Vince, A. Enumeration of Viral Capsid Assembly Pathways: Tree Orbits Under Permutation Group Action. *Bull. Math. Biol.* **2011**, *73*, 726–753.

- (29) Sitharam, M.; Ozkan, A.; Pence, J.; Peters, J. EASAL: Efficient Atlasing, Analysis and Search of Molecular Assembly Landscapes. 2004, arXiv:1203.3811. arXiv.org e-Print archive. <http://arxiv.org/abs/1203.3811> (accessed May 22, 2016).
- (30) Zlotnick, A. Theoretical Aspects of Virus Capsid Assembly. *J. Mol. Recognit.* **2005**, *18*, 479–490.
- (31) Zlotnick, A. Distinguishing Reversible from Irreversible Virus Capsid Assembly. *J. Mol. Biol.* **2006**, *366*, 14–18.
- (32) Burns, K.; Mukherjee, S.; Keef, T.; Johnson, J. M.; Zlotnick, A. Altering the Energy Landscape of Virus Self-assembly to Generate Kinetically Trapped Nanoparticles. *Biomacromolecules* **2009**, *11*, 439–442.
- (33) Lei, H.; Yang, X.; Zheng, B.; Lin, G.; Baker, N. A. Quantifying the Influence of Conformational Uncertainty in Biomolecular Solvation. 2014, arXiv:1408.5629. arXiv.org e-Print archive. <http://arxiv.org/abs/1408.5629> (accessed May 22, 2016).
- (34) Rasheed, M.; Clement, N.; Bhowmick, A.; Bajaj, C. Quantifying and Visualizing Uncertainties in Molecular Models. arXiv.org e-Print archive. <http://arxiv.org/abs/1508.03882v2>, 2015.
- (35) Hoeffding, W. Probability Inequalities for Sums of Bounded Random Variables. *J. Am. Stat. Assoc.* **1963**, *58*, 13–30.
- (36) Azuma, K. Weighted Sums of Certain Dependent Random Variables. *Tokoku Mathematical Journal* **1967**, *19*, 357–367.
- (37) McDiarmid, C. On the Method of Bounded Differences. *Surveys in Combinatorics* **1989**, *141*, 148–188.
- (38) Niederreiter, H. Quasi-Monte Carlo Methods. *Encyclopedia of Quantitative Finance* **1990**, *24*, 55–61.

- (39) Hoogland, J.; James, F.; Kleiss, R. In *Monte Carlo and Quasi-Monte Carlo Methods 1996: Proceedings of a conference at the University of Salzburg, Austria, July 9–12, 1996*; Niederreiter, H., Hellekalek, P., Larcher, G., Zinterhof, P., Eds.; Springer New York: New York, NY, 1998; Chapter Quasi-Monte Carlo, Discrepancies and Error Estimates, pp 266–276.
- (40) Bajaj, C.; Bhowmick, A.; Chattopadhyay, E.; Zuckerman, D. On Low Discrepancy Samplings in Product Spaces of Motion Groups. arXiv.org e-Print archive. <http://arxiv.org/abs/1411.7753>, 2014.
- (41) Wilcoxon, F. Individual Comparisons by Ranking Methods. *Biom. Bull.* **1945**, *1*, 80–83.
- (42) Emekli, U.; Schneidman-Duhovny, D.; Wolfson, H. J.; Nussinov, R.; Haliloglu, T. HingeProt: Automated Prediction of Hinges in Protein Structures. *Proteins: Struct., Funct., Bioinf.* **2008**, *70*, 1219–1227.
- (43) Niederreiter, H. *Random Number Generation and Quasi-Monte Carlo Methods*; Society for Industrial and Applied Mathematics: Philadelphia, PA, USA, 1992.
- (44) Gopalan, P.; Meka, R.; Reingold, O.; Zuckerman, D. Pseudorandom Generators for Combinatorial Shapes. *SIAM J. Comput.* **2013**, *42*, 1051–1076.
- (45) Bajaj, C.; Chowdhury, R. A.; Rasheed, M. A Dynamic Data Structure for Flexible Molecular Maintenance and Informatics. *Bioinformatics* **2011**, *27*, 55–62.
- (46) Bajaj, C.; Zhao, W. Fast Molecular Solvation Energetics and Forces Computation. *SIAM Journal on Scientific Computing* **2010**, *31*, 4524–4552.
- (47) Cha, D.; Zhang, Q.; Tithi, J. J.; Rand, A.; Chowdhury, R. A.; Bajaj, C. Accelerated Molecular Mechanical and Solvation Energetics on Multicore CPUs and Manycore GPUs. Proceedings of the 6th ACM Conference on Bioinformatics, Computational Biology and Health Informatics. 2015; pp 222–231.

- (48) Case, D. A.; Cheatham, T. E.; Darden, T.; Gohlke, H.; Luo, R.; Merz, K. M.; Onufriev, A.; Simmerling, C.; Wang, B.; Woods, R. J. The Amber Biomolecular Simulation Programs. *J. Comput. Chem.* **2005**, 26, 1668–1688.
- (49) Zlotnick, A. Theoretical Aspects of Virus Capsid Assembly. *J. Mol. Recognit.* **2005**, 18, 479–490.
- (50) Yedidia, J. S.; Freeman, W. T.; Weiss, Y. Understanding belief propagation and its generalizations. *Exploring artificial intelligence in the new millennium* **2003**, 8, 236–239.
- (51) Kschischang, F. R.; Frey, B. J.; Loeliger, H.-A. Factor graphs and the sum-product algorithm. *Information Theory, IEEE Transactions on* **2001**, 47, 498–519.